



SUPERCONDUCTING TECHNOLOGY ASSESSMENT



National Security Agency
Office of Corporate Assessments

AUGUST 2005

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE AUG 2005		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Superconducting Technology Assessment				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NSA Office of Corporate Assessments 9800 Savage Road, Suite 6468 Fort George G. Meade, Maryland 20755				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 257	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

August 2005

Superconducting Technology Assessment Letter of Promulgation

Background

This Superconducting Technology Assessment (STA) has been conducted by the National Security Agency to address the fundamental question of a potential replacement for silicon complementary metal oxide semiconductor (CMOS) in very high-end computing (HEC) environments. Recent industry trends clearly establish that design tradeoffs between power, clock and metrology have brought CMOS to the limits of its scalability. All microprocessor firms have turned to multiple cores and reduced power in efforts to improve performance. Increased parallelism on a chip permits some architectural innovation, but it also increasingly shifts issues of performance gains into software application environments, where there are already many practical limits to scalability of performance. For many demanding applications in the U. S. national security, scientific, medical and industrial sectors, availability of higher-performance components in well-balanced HEC environments is essential. Alternatives to CMOS must therefore be found.

The Semiconductor Industry Association (SIA) International Technology Roadmap for Semiconductors (ITRS) has identified Superconducting Rapid Single Flux Quantum (RSFQ) technology as the most promising technology in the continuing demand for faster processors. There has been steady progress in research in this technology, though with somewhat weaker efforts at development and industrialization. This assessment is an in-depth examination of RSFQ technologies with the singular objective of determining if a comprehensive roadmap for technology development is possible, aiming for industrial maturity in the 2010-2012 timeframe. The goal would be an RSFQ technology set sufficient to support development of true petaflop-scale computing at the end of this decade.

Methodology

A team of national experts in superconducting technologies was empanelled to conduct this assessment. It was chaired by Dr. John Pinkston, former Chief of Research, NSA. Membership is shown in Appendix B and included experts in processor architectures, several types of memories, interconnects, design and manufacturing in these technologies. The panel heard from academic, industrial and government experts in the field and reviewed available superconducting documentation. The panel also had the benefit of presentations and discussions with two leading HEC architects on system-level issues that could conceivably impact the use of RSFQ technologies. The panel had lengthy internal discussions on the mutual dependencies of various superconducting components. Availability, technical issues, development schedules and potential costs were discussed at length. The resulting roadmap represents a consensus of views with no substantial dissension among panel members. The panel was enjoined from examining HEC architectural issues and systems-level options, other than those bearing directly on technology envelopes (e.g., is a 50 Ghz clock a sufficient goal?).

Summary of Findings

The STA concluded that there were no significant outstanding research issues for RSFQ technologies. Speed, power and Josephson Junction density projections could be made reliably. Areas of risk have been identified and appropriately dealt with in the roadmap, with cautionary comments on mitigation or alternatives. Memories are clearly one such area, but this report concludes that the field suffers more from lack of research than available alternatives. The assessment, in fact, identifies several memory alternatives, each of which should be pursued until technology limits are clearly understood. Development of RSFQ technologies to sufficient levels of maturity, with appropriate milestones, could be achieved in the time frame of interest but would require a comprehensive and sustained government funded program of approximately \$100M/yr. The panel concluded that private sector interests in superconducting RSFQ would not be sufficient to spawn development and industrialization.

It is, of course, NSA's role to turn these STA findings into specific actions in partnership with national security community and other federal HEC users having extreme high end computing needs and the vision to pursue the performance goals that superconducting RSFQ appears to offer.

The undersigned government members of this study would like to commend our industry and academia participants for the balanced and constructive assessment results.




Dr. Fernand Bedard



Dr. Nancy K. Welker



George R. Cotter



Michael A. Escavage



Dr. John T. Pinkston

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Assessment Objective and Findings	01
Limitations of Current Technology	02
Panel Tasked	02
RSFQ Technology Is Viable	03
Roadmap Created	04
Government Investment Crucial	05
Technical Issues	05
Conclusions	05

CHAPTER 01: INTRODUCTION

1.1 NSA Dependent on High-End Computing	07
1.2 Limitations of Conventional Technology for High-End Computing	08
1.2.1 Conventional Silicon Technology Not the Answer	08
1.2.2 Supercomputing RSFQ a Viable Alternative	08
1.3 What is RSFQ Circuitry?	09
1.3.1 Josephson Junctions	09
1.3.2 RSFQ Attributes	10
1.4 Summary of Panel's Efforts	10
1.4.1 RSFQ Ready for Investment	11
1.4.2 RSFQ Can Leverage Microprocessor Technology	11
1.5 Roadmap Created and Government Investment Needed	11
1.5.1 Funding	13
1.6 Technical Issues	13
1.6.1 Supercomputing System Constraints on the Use of RSFQ	13
1.6.2 Fast, Low Latency Memory	15
1.6.3 High-Speed Input/Output	15
1.6.4 CAD Tools	15
1.6.5 Refrigeration	15
1.7 State of the Industry	15
1.8 Contents of Study	16
1.9 Chapter Summaries	16

CHAPTER 02: ARCHITECTURAL CONSIDERATIONS FOR SUPERCONDUCTOR RSFQ MICROPROCESSORS

2.1	Superconductor Microprocessors—Opportunities, Challenges, and Projections	25
2.2	Microprocessors – Current Status of RSFQ Microprocessor Design	28
2.2.1	SPELL Processors for the HTMT Petaflops System (1997-1999)	28
2.2.2	20-GHz, 8-bit FLUX-1 Microprocessor (2000-2002)	29
2.2.3	CORE1 Bit-serial Microprocessor Prototypes (2002-2005)	30
2.2.4	Proposal for an RSFQ Petaflops Computer in Japan (est. 2005-2015)	30
2.3	Microprocessors – Readiness	32
2.4	Microprocessors – Issues and Concerns	33
2.4.1	Clocking for 50 GHz RSFQ Processors	33
2.4.2	Long Processing Pipelines	33
2.4.3	On-chip Interconnect, Chip Area Reachable in a Single Cycle, and Register File Structures	33
2.4.4	Memory Hierarchy for Superconductor Processors and Systems	34
2.4.5	Memory Latency Tolerant Architectures	35
2.5	Microprocessors – Conclusions and Goals	36
2.6	Microprocessors – Roadmap and Milestones	37
2.7	Microprocessors – Funding	38

CHAPTER 03: SUPERCONDUCTIVE RSFQ PROCESSOR AND MEMORY TECHNOLOGY

3.1	RSFQ Processors	43
3.1.1	RSFQ Processors – Status	43
3.1.2	RSFQ Processors – Readiness for Major Investment	47
3.1.3	RSFQ Processors – Roadmap	49
3.1.4	RSFQ Processors – Investment Required	50
3.1.5	RSFQ Processors – Issues and Concerns	51
3.1.6	RSFQ Processors – Projected Future Capability	52
3.2	Memory	52
3.2.1	Memory – Hybrid Josephson-CMOS RAM	54
3.2.2	Memory – Single-Flux Quantum Memory	56
3.2.3	Memory – MRAM	59
3.2.4	Memory – Summation	66
3.3	CAD Tools and Design Methodologies	67
3.4	Summary	72

CHAPTER 04: SUPERCONDUCTIVE CHIP MANUFACTURE

4.1	SCE IC Chip Manufacturing – Scope	78
4.2	Digital Superconductive IC Chip Fabrication – Status	79
4.3	SCE Chip Fabrication for HEC – Readiness	82
4.4	IC Chip Manufacturing Capability – Projections	83
4.5	IC Chip Manufacture – Issues and Concerns	84
4.5.1	IC Chip Manufacture – IC Manufacturability	84
4.5.2	IC Chip Manufacture – Device and Circuit Speed	85
4.5.3	IC Chip Manufacture – Circuit Density and Clock Speed	86
4.5.4	IC Chip Manufacture – Parameter Spreads and IC Chip Yield	87
4.5.5	IC Chip Manufacture – Production of Known Good Die	89
4.6	Roadmap and Facilities Strategy	90
4.6.1	Roadmap and Facilities Strategy – Roadmap	90
4.6.2	Roadmap and Facilities Strategy – Manufacturing Facilities	92

CHAPTER 05: INTERCONNECTS AND SYSTEM INPUT/OUTPUT

5.1	Optical Interconnect Technology	99
5.1.1	Optical Interconnect Technology – Status	100
5.1.2	Optical Interconnect Technology – Readiness	101
5.1.3	Optical Interconnect Technology – Projections	102
5.1.4	Optical Interconnect Technology – Issues and Concerns:	103
5.2	Input: Data and Signal Transmission from Room Temperature to 4 K	103
5.2.1	Input: Room Temperature to 4 K – Status	104
5.2.2	Input: Room Temperature to 4 K – Issues and Concerns	104
5.2.3	Input: Room Temperature to 4 K – Roadmap	104
5.3	Output: 4 K RSFQ to Room Temperature Electronics	105
5.3.1	Output: 4 K RSFQ to Room Temperature Electronics – Status	105
5.3.2	Output: 4 K RSFQ to Room Temperature Electronics – Readiness and Projections	106
5.3.3	Output: 4 K RSFQ to Room Temperature Electronics – Issues and Concerns	107
5.3.4	Output: 4 K RSFQ to Room Temperature Electronics – Roadmap and Funding Profile	108
5.4	Data Routing: 4 K RSFQ to 4 K RSFQ	108
5.4.1	Data Routing: 4 K RSFQ to 4 K RSFQ – Status	109
5.4.2	Data Routing: 4 K RSFQ to 4 K RSFQ – Readiness	110
5.4.3	Data Routing: 4 K RSFQ to 4 K RSFQ – Issues and Concerns	110
5.4.4	Data Routing: 4 K RSFQ to 4 K RSFQ – Roadmap and Funding	111

CHAPTER 06: SYSTEM INTEGRATION

6.1	Multi-chip Modules and Boards	114
6.1.1	Multi-chip Modules and Boards – Status	114
6.1.2	Multi-chip Modules and Boards – Readiness	115
6.1.3	Multi-chip Modules and Boards – Projections	116
6.1.4	Multi-chip Modules and Boards – Issues and Concerns	116
6.1.5	Multi-chip Modules and Boards – Roadmap and Funding	117
6.2	3-D Packaging	118
6.2.1	3-D Packaging – Status	118
6.2.2	3-D Packaging – Readiness	119
6.2.3	3-D Packaging – Issues and Concerns	120
6.2.4	3-D Packaging – Roadmap and Funding	120
6.3	Enclosures and Shields	121
6.3.1	Enclosures and Shields – Status	121
6.3.2	Enclosures and Shields – Readiness	122
6.3.3	Enclosures and Shields – Projections	122
6.3.4	Enclosures and Shields – Issues and Concerns	122
6.3.5	Enclosures and Shields – Roadmap and Funding	123
6.4	Cooling	123
6.4.1	Cooling – Status	124
6.4.2	Cooling – Readiness	125
6.4.3	Cooling – Issues and Concerns	125
6.4.4	Cooling – Roadmap and Funding	126
6.5	Power Distribution and Cables	127
6.5.1	Power Distribution and Cables – Status	127
6.5.2	Power Distribution and Cables – Readiness	129
6.5.3	Power Distribution and Cables – Projections	129
6.5.4	Power Distribution and Cables – Issues and Concerns	129
6.5.5	Power Distribution and Cables – Roadmap and Funding	130
6.6	System Integrity and Testing	131
6.6.1	System Integrity and Testing – Status	131
6.6.2	System Integrity and Testing – Readiness	131
6.6.3	System Integrity and Testing – Issues and Concerns	131
6.6.4	System Integrity and Testing – Roadmap and Funding	132

APPENDIX A: TERMS OF REFERENCE	135
APPENDIX B: PANEL MEMBERS	139
APPENDIX C: GLOSSARY	143
APPENDIX D: INTRODUCTION TO SUPERCONDUCTOR SINGLE FLUX QUANTUM CIRCUITRY	147
 (The appendices below can be found on the attached CD but are not printed in this report.)	
APPENDIX E: SOME APPLICATIONS FOR RSFQ	155
APPENDIX F: SYSTEM ARCHITECTURES	161
APPENDIX G: ISSUES AFFECTING RSFQ CIRCUITS	169
APPENDIX H: MRAM TECHNOLOGY FOR RSFQ HIGH-END COMPUTING	177
APPENDIX I: SUPERCONDUCTOR INTEGRATED CIRCUIT FABRICATION TECHNOLOGY	185
APPENDIX J: CAD	203
APPENDIX K: DATA SIGNAL TRANSMISSION	213
APPENDIX L: MULTI-CHIP MODULES AND BOARDS	227

INDEX OF FIGURES

CHAPTER 01: INTRODUCTION

1-1	2004 ITRS update shows RSFQ as the lowest risk potential emerging technology for processing beyond silicon	09
1-2	Roadmap for RSFQ technology tools and components	12
1-3	Conceptual architecture	14

CHAPTER 02: ARCHITECTURAL CONSIDERATIONS FOR SUPERCONDUCTOR RSFQ MICROPROCESSORS

2-1	8-processor node of the proposed Japanese superconductor petaflops system	31
2-2	50-GHz multi-chip processor microprocessor design roadmap	37

CHAPTER 03: SUPERCONDUCTIVE RSFQ PROCESSOR AND MEMORY TECHNOLOGY

3.1-1	Roadmap and major milestones for RSFQ processors	43
3.1-2	Photograph of 1-cm ² Flux-1 63K JJ microprocessor chip	44
3.1-3	Signal and ground pad geometry used on Flux-1	47
3.2-1	Hybrid Josephson-CMOS RAM operates at 4 K	54
3.2-2	Wafer with embedded CMOS memory chips for direct wiring of Josephson peripheral circuits	55
3.2-3	Diagram of FS-TMR cell	61
3.2-4	Cryogenic RAM roadmap	66
3.3-1	ADC chip containing about 6,000 JJs	68
3.3-2	A bit-serial superconductor microprocessor featuring 6,300 JJs, 7 instructions, and a 16 GHz clock	68

CHAPTER 04: SUPERCONDUCTIVE CHIP MANUFACTURE

4-1	Schematic diagram of functions included in IC chip manufacture	78
4-2	Fabrication of RSFQ ICs is accomplished using semiconductor equipment and processes	79
4-3	Demonstrations of RSFQ circuit speed with increasing J _c	80
4-4	Notional cross-section of the superconductive IC chip fabrication process illustrating salient features of an advanced process which includes four interconnect levels, planarization, and plugs	85
4-5	Projections of RSFQ circuit speeds with increasing J _c	87
4-6	Timeline for development of SCE manufacturing capability	92

CHAPTER 05: INTERCONNECTS AND SYSTEM INPUT/OUTPUT

5-1	A 64-fiber, 4-wavelength, 25-Gbps CWDM System for bi-directional transmission totaling 6.4 Tbps between a superconducting processor at 4 K and high speed mass memory at 300 K	100
5-2	Four channel transceiver arrays operating at 10 Gbps/channel produced by IBM/Agilent	101
5-3	The use of crossbar switching in supercomputers	109
5-4	Superconducting 16x16 crossbar switch with 14,000 junctions on a 5x5 mm chip	110

CHAPTER 06: SYSTEM INTEGRATION

6-1a	System installation concept for petaflops HTMT system	113
6-1b	Packaging concept for HTMT SCP	113
6-2	HTMT conceptual packaging for cryogenic processing and data communications	114
6-3	Multi-chip module with SCE chips	115
6-4	Categorization of 3D packaging approaches	119
6-5	Construction details of a typical enclosure designed for operation at 4 K	121
6-6	Concept for a large-scale system including cryogenic cooling unit for supercomputers	125
6-7	High-speed flexible ribbon cable designed for modular attachment	128

APPENDIX A: TERMS OF REFERENCE (none)

APPENDIX B: PANEL MEMBERS (none)

APPENDIX C: GLOSSARY (none)

APPENDIX D: INTRODUCTION TO SUPERCONDUCTOR SINGLE FLUX QUANTUM CIRCUITRY

Figure 1:	Equivalent circuit of Josephson junction. I_C represents the nonlinear switch	147
Figure 2a:	DC electrical characteristics of voltage-state latching junctions	148
Figure 2b:	DC electrical characteristics of SFQ non-latching junctions	148
Figure 3:	Measured speed of static dividers varies as $J_C^{1/2}$	149
Figure 4:	Passive transmission lines can propagate picosecond SFQ pulses without dispersion at the speed of light in the line	150
Figure 5:	Representative SFQ gates	151

(The appendices below can be found on the attached CD but are not printed in this report.)

APPENDIX E: SOME APPLICATIONS FOR RSFQ (none)

APPENDIX F: SYSTEM ARCHITECTURES (none)

APPENDIX G: ISSUES AFFECTING RSFQ CIRCUITS

Figure 1: Calculated BER of a 2-JJ comparator as a function of operating margin	169
Figure 2: Yield as a function of process variations	171

APPENDIX H: MRAM TECHNOLOGY FOR RSFQ HIGH-END COMPUTING

Figure 1: Schematic of a 1T-1MTJ MRAM cell structure showing the sense path and programming lines	177
Figure 2: Schematic of a proposed SMT MRAM cell structure showing the common current path for sense and program operations	178
Figure 3: Estimated cell sizes for the two MRAM technologies compared to NOR Flash for IC technology nodes from 90nm to 32nm	179
Figure 4: Estimated switching currents for SMT devices at various lithography nodes	182

APPENDIX I: SUPERCONDUCTOR INTEGRATED CIRCUIT FABRICATION TECHNOLOGY

Figure 1: Cross section of NGST's 8-kA/cm ² niobium-based 14-mask step integrated circuit process	1519
Figure 2: Current density versus oxidation pressure-time product for J _c of 2-20 kA/cm ²	1521
Figure 3: Trend chart of current density	1521
Figure 4: Junction fabrication process showing key features	1522
Figure 5: SEM photograph of a 1.0 μm junction and self-aligned anodization ring on base electrode	1523
Figure 6: SEM photograph of a partially completed T-flip flop stage showing junction, anodization ring and base electrode	1523
Figure 7: Typical current-voltage characteristics of a 8-kA/cm ² 1.25-μm diameter junction	1523
Figure 8: Typical current-voltage characteristics of a series array of one hundred 8-kA/cm ² 1.25-μm diameter junctions	1523
Figure 9: SEM photograph showing the reentrant step coverage of sputter-deposited SiO ₂ over a 500-nm metal step	1524
Figure 10: SEM photograph showing the improvement in step coverage of SiO ₂ over a 500-nm metal step with substrate bias applied during deposition	1525
Figure 11: Ground plane planarization process	1526
Figure 12: Trend chart of leakage conductance for comb-to-meander structure showing that oxide fill eliminates shorts between adjacent comb-to-meander wires over ground etch meander in PCM test structure	1527
Figure 13: SEM picture of first and second wiring layers crossing ground moat filled with oxide	1527
Figure 14: Maximum reported TRFF divider speed	1528
Figure 15: Typical current-voltage characteristics of a 20-kA/cm ² , 0.90-μm junction	1529
Figure 16: Plot and linear fit of the square-root-I _c versus drawn junction diameter for a series of 20-kA/cm ² junctions ranging in diameter	1529
Figure 17: Photograph of the FLUX-1r1 microprocessor chip fabricated in NGST's 4-kA/cm ² process	1530

APPENDIX J: CAD

Figure 1: Circuit design and verification	203
Figure 2: Schematic view	204
Figure 3: Symbol view	205
Figure 4: VHDL view	206
Figure 5: Physical layout view	207
Figure 6: LMeter is specialized software that extracts the inductance value of an interconnect from the physical layout	208
Figure 7: VHDL simulation performed on a large, mixed signal superconductor circuit	209
Figure 8: Layout-versus-Schematic verification on the chip level for a large chip	210
Figure 9: First-pass success is routine for circuits of up to a few thousand JJs	211

APPENDIX K: DATA SIGNAL TRANSMISSION

Figure 1: Channel loss for 50 cm electrical link	214
Figure 2: Long path data signal transmission requirements for one proposes petaflop architecture	215
Figure 3: Four channel transceiver arrays produced by the IBM/Agilent effort funded by DARPA	216
Figure 4: A 64-fiber, 4-wavelength, 25-Gbps CWDM system for bidirectional transmission totaling 6.4 Tbps between a superconductive processor at 4 K and high speed mass memory at 300K	220
Figure 5: A 3-fiber, 64-wavelength, 50-Gbps DWDM system for bidirectional transmission totaling 6.4 Tbps between a superconductive processor at 4 K and high speed mass memory at 300K	221

APPENDIX L: MULTI-CHIP MODULES AND BOARDS

Figure 1: Packaging concept for HTMT SCP	227
Figure 2: A multi-chip module with SCE chips	228
Figure 3: Lead inductance vs. bond length or height for different attach techniques	229
Figure 4: Heterogeneous stacking technology for system-in-stack	232
Figure 5: Stacks containing superconducting electronics circuits were cycled from RT to 4 K several times	232
Figure 6: A prototype cryo enclosure for operation at 4 K	233
Figure 7: A typical cryocooler enclosure designed for communication applications	234
Figure 8: A rack mounted cryocooler enclosure example	234
Figure 8: Northrop Grumman's high-efficiency cryocooler unit	236
Figure 9: Refrigerator efficiency of various cooling systems	237
Figure 10: Concept for a large-scale system including cryogenic cooling unit for supercomputers	237
Figure 11: The cost of various refrigerators as a function of refrigeration	239
Figure 12: A high-speed flexible ribbon cable designed for modular attachment	240

CONTENTS OF CD

EXECUTIVE SUMMARY

CHAPTER 01: INTRODUCTION

CHAPTER 02: ARCHITECTURAL CONSIDERATIONS FOR SUPERCONDUCTOR RSFQ MICROPROCESSORS

CHAPTER 03: SUPERCONDUCTIVE RSFQ PROCESSOR AND MEMORY TECHNOLOGY

CHAPTER 04: SUPERCONDUCTIVE CHIP MANUFACTURE

CHAPTER 05: INTERCONNECTS AND SYSTEM INPUT/OUTPUT

CHAPTER 06: SYSTEM INTEGRATION

APPENDIX A: TERMS OF REFERENCE

APPENDIX B: PANEL MEMBERS

APPENDIX C: GLOSSARY

APPENDIX D: INTRODUCTION TO SUPERCONDUCTOR SINGLE FLUX QUANTUM CIRCUITRY

APPENDIX E: SOME APPLICATIONS FOR RSFQ

APPENDIX F: SYSTEM ARCHITECTURES

APPENDIX G: ISSUES AFFECTING RSFQ CIRCUITS

APPENDIX H: MRAM TECHNOLOGY FOR RSFQ HIGH-END COMPUTING

APPENDIX I: SUPERCONDUCTOR INTEGRATED CIRCUIT FABRICATION TECHNOLOGY

APPENDIX J: CAD

APPENDIX K: DATA SIGNAL TRANSMISSION

APPENDIX L: MULTI-CHIP MODULES AND BOARDS

All financial figures throughout the report are listed in millions of dollars (\$M); all financial figures in chapter summaries are rounded to the nearest million.

EXECUTIVE SUMMARY

ASSESSMENT OBJECTIVE AND FINDINGS

The government, and particularly NSA, has a continuing need for ever-increasing computational power. The Agency is concerned about projected limitations of conventional silicon-based technology and is searching for possible alternatives to meet its future mission-critical computational needs.

This document presents the results of a Technology Assessment, chartered by the Director of NSA, to assess the readiness of ultra-high-speed superconductive (SC) Rapid Single Flux Quantum (RSFQ) circuit technology for application to very-high-performance (petaflops-scale) computing systems. A panel of experts performed this assessment and concluded that:

- RSFQ technology is an excellent candidate for petaflops-scale computers.
- Government investment is necessary, because private industry currently has no compelling financial reason to develop this technology for mainstream commercial applications.
- With aggressive federal investment (estimated between \$372 and \$437 million over five years), by 2010 RSFQ technology can be sufficiently matured to allow the initiation of the design and construction of an operational petaflops-scale system.
- Although significant risk issues exist, the panel has developed a roadmap that identifies the needed technology developments with milestones and demonstration vehicles.

TABLE E-1. REASONS TO DEVELOP SUPERCONDUCTIVE COMPUTER TECHNOLOGY	
Technological	Financial
NSA’s computing needs are outstripping conventional technology.	Market forces alone will not drive private industry to develop SC technology.
RSFQ technology is an excellent candidate for higher-performance computing capability.	The federal government will be the primary end user of SC computer technology.
RSFQ technology has a clear and viable roadmap.	Other federal government missions will benefit from advances in SC technology.

LIMITATIONS OF CURRENT TECHNOLOGY

Circuit Speeds Facing Limits

The government is observing increased difficulty as industry attempts to raise the processing performance of today's silicon-based supercomputer systems through improvements in circuit speeds. In the past several decades, steady decreases in circuit feature sizes have translated into faster speeds and higher circuit densities that have enabled ever-increasing performance. However, conventional technology has a limited remaining lifetime and is facing increasing challenges in material science and power dissipation at smaller feature sizes.

There are already signs that the major commodity device industry is turning in other directions. The major microprocessor companies are backing away from faster clock speeds and instead are fielding devices with multiple processor "cores" on a single chip, with increased performance coming from architectural enhancements and device parallelism rather than increased clock speed.

While the Semiconductor Industry Association (SIA) International Technology Roadmap for Semiconductors (ITRS) projects silicon advances well into the next decade, large-scale digital processing improvements will almost certainly come from increased parallelism, not raw speed.

Commercial and Government Interests Diverging

Over the past two decades, High-End Computing (HEC) systems have improved by leveraging a large commodity microprocessor and consumer electronics base. However, future evolution of this base is projected to diverge from the technology needs of HEC for national security applications by supporting more processors rather than faster ones. The result will be limitations in architecture and programmability, for implementations of HEC based on the traditional commodity technology base.

Power Requirements Swelling

For supercomputers, continuing reliance on this technology base means a continuation of the trend to massively parallel systems with thousands of processors. However, at today's scale, the electrical power and cooling requirements are bumping up against practical limits, even if ways were found to efficiently exploit the parallelism. For example, the Japanese Earth Simulator system, which has been ranked number one on the list of the top 500 installed HEC, consumes approximately 6 megawatts of electrical power.

PANEL TASKED

A panel of experts from industry and academia, augmented by Agency subject matter experts, was assembled to perform this study, bringing expertise from superconducting materials, circuitry, fabrication, high-performance computer architecture, optical communications, and other related technologies. The panel:

- Assessed RSFQ technology for application to high-performance computing systems available after 2010, based on current projections of material science, device technology, circuit design, manufacturability, and expected commercial availability of superconductive (SC) technology over the balance of the decade.
- Identified roadmaps for the development of the essential components, including microprocessor circuits, memory, and interconnect, for high-end computer architectures by 2010.

RSFQ TECHNOLOGY IS VIABLE

TABLE E-2. RSFQ SUMMARY	
Technical Advantages	Technical Challenges
The most advanced alternative technology.	Providing high-speed and low-latency memory.
Combines high speed with low power.	Architecting systems that can tolerate significant memory access latencies.
Ready for aggressive investment.	Providing very high data rate communications between room temperature technology and cooled RSFQ.

Most Advanced Alternative Technology

The ITRS 2004 update on Emerging Research Devices lists many candidate technologies, presently in the research laboratories, for extending performance beyond today's semiconductor technology. Superconducting RSFQ is included in this list and is assessed to be at the most advanced state of any of the alternative technologies.

Ready for Aggressive Investment

In the opinion of the panel, superconducting RSFQ circuit technology is ready for an aggressive, focused investment to meet a 2010 schedule for initiating the development of a petaflops-class computer. This judgment is based on:

- An evaluation of progress made in the last decade.
- Projection of the benefits of an advanced very-large-scale integration (VLSI) process for RSFQ in a manufacturing environment.
- A roadmap for RSFQ circuit development coordinated with VLSI manufacturing and packaging technologies.

Can Leverage Semiconductor Technology Base







Although RSFQ circuits are still relatively immature, their similarity in function, design, and fabrication to semiconductor circuits permits realistic extrapolations. Most of the tools for design, test, and fabrication are derived directly from the semiconductor industry, although RSFQ technology will still need to modify them. Significant progress has already been demonstrated on limited budgets by companies such as Northrop Grumman and HYPRES, and in universities such as Stony Brook and the University of California, Berkeley.

High Speed with Low Power

Individual RSFQ circuits have been demonstrated operating at clock rates in the hundreds of GHz, and system clocks of at least 50GHz seem quite attainable, with faster speeds possible. In addition, RSFQ devices have lower power requirements than other systems, even after cooling requirements are included. Extremely low RSFQ power enables compact systems with greatly increased computational capability for future government needs, but with no increase in overall power requirements beyond today's high-end systems.

State of the Industry

Today, expertise in digital SC technology resides in only a handful of companies and institutes. The U.S. base is shrinking:

TABLE E-3. DIGITAL RSFQ TECHNOLOGY'S CURRENT STATE OF THE INDUSTRY		
Country	Entity	Status
	ISTEC/SRL	<ul style="list-style-type: none"> – Joint government/industry center, probably doing the most advanced work in digital RSFQ anywhere in the world today. – Responsible for the Earth Simulator system.
	HYPRES	<ul style="list-style-type: none"> – Private company focused entirely on SC digital electronics. – Has operated the only full-service commercial foundry in the U.S. since 1983.
	Northrop Grumman	<ul style="list-style-type: none"> – Had the most advanced foundry and associated design capability until suspended last year. – Still has a strong cadre of experts in the field.
	Stony Brook U, UC Berkeley, JPL	<ul style="list-style-type: none"> – Currently conducting academic research.
	Chalmers U of Technology	<ul style="list-style-type: none"> – Currently conducting academic research.
	NSA, NIST	<ul style="list-style-type: none"> – Have resident expertise.

ROADMAP CREATED

This report presents a detailed RSFQ technology roadmap, defining the tools and components essential to support RSFQ-based high-end computing by 2010. The end point of this roadmap includes:

- An RSFQ processor of approximately 1-million gate complexity, operating at a 50 GHz clock rate.
- A design capability:
 - consisting of an RSFQ cell library and a complete suite of computer-aided design (CAD) tools.
 - allowing a competent ASIC digital designer with no background in superconductor electronics to design high-performance processor units.
- An RSFQ chip manufacturing facility with an established stable process operating at high yields.

GOVERNMENT INVESTMENT CRUCIAL

There is no foreseeable commercial demand for SC digital technology products sufficient to justify significant private industry investment in developing that technology. For this reason, government funding is crucial to this technology's development. Besides its use to NSA, SC will likely have applications for other government missions as well. Once it has been sufficiently developed, SC may also prove to have commercial applications.

While the panel finds RSFQ technology very promising as a base for future HEC systems, this technology will still require significant developmental effort and an investment of between \$372 and \$437 million over five years in order to be ready for design and construction of operational systems.

TECHNICAL ISSUES

Many technical problems remain to be solved on the path to maturing RSFQ technology. This report presents sequences of experiments and developmental steps that would address the major issues critical to the success of this effort. Those issues, which must be addressed aggressively by any developmental program, are:

- Providing high-speed/low-latency memory.
- Architecting systems that can tolerate significant memory access latencies.
- Providing very-high-data-rate communications into and (particularly) out of the cryogenic environment.

Other technical issues—considered to be reasonably low risk, although still in need of development work—include:

- Providing powerful CAD tools for the designers.
- Achieving a stable fabrication process.
- Refrigeration.

CONCLUSIONS

RSFQ technology is ready for a major development program culminating in demonstration vehicles that will open the door to operational systems. This can be accomplished in five years with an aggressive, government funded program. Without such government investment, this development will not happen.

01

This document presents the results of a Technology Assessment, chartered by the Director of NSA, to assess the readiness of an ultra-high-speed circuit technology, superconductive Rapid Single Flux Quantum (RSFQ), for use in very-high-performance (petaflops-scale) computing systems.

INTRODUCTION

The request for this assessment was motivated by the government's assessment that conventional technology for high-end computing (HEC) systems is finding it more and more difficult to achieve further increases in computational performance. A panel of experts in superconductive electronics, high-performance computer architectures, and related technologies was formed to conduct this study. The composition of the panel is presented in Appendix B.

In summary, the charge to the panel was to:

"...conduct an assessment of superconducting technology as a significant follow-on to silicon for component use in high performance computing systems available after 2010. The assessment will examine current projections of material science, device technology, circuit design, manufacturability, and general commercial availability of superconducting technology over the balance of the decade, identify programs in place or needed to advance commercialization of superconducting technology if warranted by technology projections, and identify strategic partnerships essential to the foregoing. First order estimates of the cost and complexity of government intervention in technology evolution will be needed. The assessment will not directly investigate potential high-end computer architectures or related non-superconducting technologies required by high-end computers other than those elements essential to the superconducting technology projections."

The full text of the formal charge to the panel is given in Appendix A.

1.1 NSA DEPENDENT ON HIGH-END COMPUTING

The NSA mission is dependent on HEC for cryptanalysis, natural language processing, feature recognition from image analysis, and other intelligence processing applications. NSA applications touch the extremes of supercomputing resource use: numerically intensive computation, high-volume data storage, and high bandwidth access to external data. As the use of electronic communications increases, so too does the volume of data to be processed and the sophistication of encryption mechanisms; thus, the need for HEC escalates.

NSA's mission is dependent on high-end computing; its applications touch the extremes of supercomputing resource use:

- Numerically intensive computation.*
 - High-volume data storage.*
 - High-bandwidth access to external data.*
-

For the past two decades, HEC systems have evolved by leveraging a large commodity microprocessor and consumer electronics base. However, there is now increased concern about the divergence of this base and the government's HEC technology needs. There are many other potential government and private sector applications of HEC. Some of these are discussed in Appendix E: Some Applications for RSFQ. (The full text of this appendix can be found on the CD accompanying this report.)

1.2 LIMITATIONS OF CONVENTIONAL TECHNOLOGY FOR HIGH-END COMPUTING

In 1999, the President's Information Technology Advisory Committee (PITAC) wrote the following in its *Report to the President-Information Technology Research: Investing in Our Future* (Feb 1999): "...Ultimately, silicon chip technology will run up against the laws of physics. We do not know exactly when this will happen, but as devices approach the size of molecules, scientists will encounter a very different set of problems fabricating faster computing components."

1.2.1 CONVENTIONAL SILICON TECHNOLOGY NOT THE ANSWER

NSA experts in HEC have concluded that semiconductor technology will not deliver the performance increases that the government's computing applications demand. Complementary metal oxide semiconductors (CMOS) is becoming less a performance technology—vendors such as Intel are voicing reluctance to seek 10 GHz clock speeds—and more a capability technology, with transistor counts of several hundred million per chip. The high transistor counts make it possible to put many functional units on a single processor chip, but then the on-chip functional units must execute efficiently in parallel. The problem becomes one of extracting parallelism from applications so that the functional units are used effectively.

Unfortunately, there are applications for which on-chip parallelism is not the solution; for such applications, blazing speed from a much smaller number of processors is required. For supercomputers, continuing reliance solely on a CMOS technology base means a continuation of the trend to massively parallel systems with thousands of processors. The result will be limitations in efficiency and programmability.

In addition, at today's scale, the electrical power and cooling requirements are facing practical limits, even if ways were found to efficiently exploit the parallelism. For example, the Japanese Earth Simulator system, which has been ranked number one on the list of the top 500 installed HEC, consumes over 6 megawatts of electrical power.

1.2.2 SUPERCOMPUTING RSFQ A VIABLE ALTERNATIVE

The Silicon Industry Association (SIA) International Technology Roadmap for Semiconductors (ITRS) 2004 update on Emerging Research Devices has many candidate technologies presently in the research laboratories for extending performance beyond today's semiconductor technology. Superconducting Rapid Single Flux Quantum (RSFQ) is included in this list and is assessed to be at the most advanced state of any of the alternative technologies. RSFQ technology has the potential to achieve circuit speeds well above 100 GHz with lower power requirements than complementary metal oxide semiconductor (CMOS), making it attractive for very-high-performance computing, as shown in Figure 1-1.

Superconducting RSFQ is assessed to be at the most advanced state of any of the alternative technologies.

Digital superconducting electronics RSFQ technology has the potential, as identified in the 2003 and 2004 SIA roadmaps, to be a “successor” technology to CMOS for high-performance applications. The 2004 Update to this roadmap stated the problem as:

“One difficult challenge related to logic in both near- and the longer-term is to extend CMOS technology to and beyond the 45 nm node sustaining the historic annual increase of intrinsic speed of high-performance MPUs at 17%. This may require an unprecedented simultaneous introduction of two or more innovations to the device structure and/or gate materials. Another longer-term challenge for logic is invention and reduction to practice of a new manufacturable information and signal processing technology addressing ‘beyond CMOS’ applications. Solutions to the first may be critically important to extension of CMOS beyond the 45 nm node, and solutions to the latter could open opportunities for microelectronics beyond the end of CMOS scaling.”

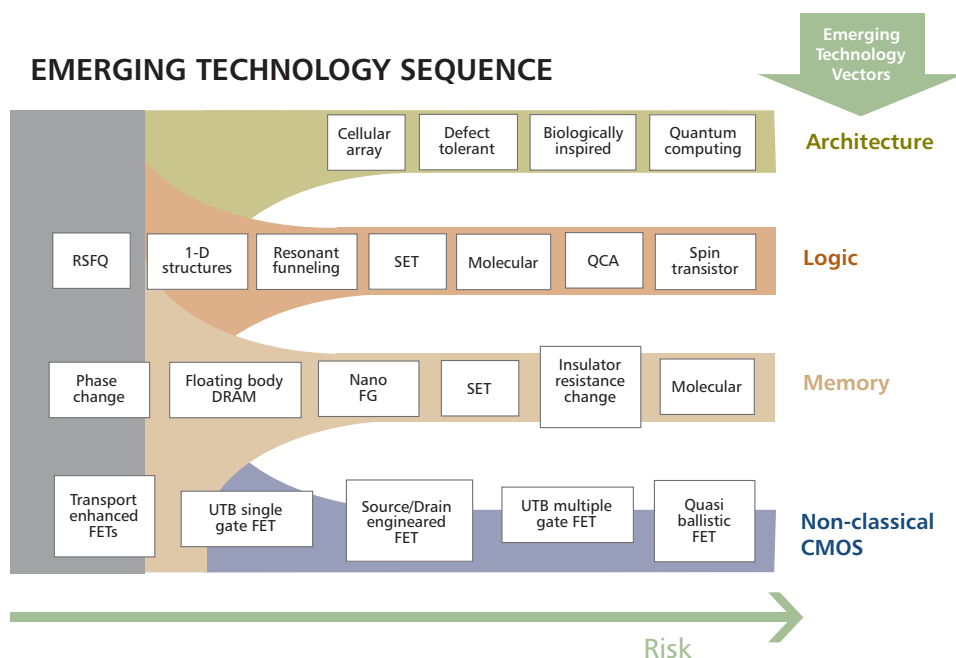


Figure 1-1. 2004 ITRS update shows RSFQ as the lowest risk (highest maturity) potential emerging technology for processing beyond silicon.

1.3 WHAT IS RSFQ CIRCUITRY?

Rapid Single Flux Quantum (RSFQ) is the latest generation of superconductor circuits based on Josephson junction devices. It uses generation, storage, and transmission of identical single-magnetic-flux-quantum pulses at rates approaching 1,000 GHz. Small asynchronous circuits have already been demonstrated at 770 GHz, and clocked RSFQ circuits are expected to exceed 100 GHz.

1.3.1 JOSEPHSON JUNCTIONS

The Josephson junction (JJ) is the basic switching device in superconductor electronics. Josephson junctions operate in two different modes: switching from zero-voltage to the voltage-state and generating single-flux quanta. The early work, exemplified by the IBM and the Japanese Josephson computer projects of the 1970's and 1980's, exclusively used logic circuits where the junctions switch between superconducting and voltage states and require AC power. RSFQ junctions generate single-flux-quantum pulses and revert to their initial superconducting condition. RSFQ circuits are DC powered.

1.3.2 RSFQ ATTRIBUTES

Important attributes of RSFQ digital circuits include:

- Fast, low-power switching devices that generate identical single-flux-quantum data pulses.
- Loss-less superconducting wiring for power distribution.
- Latches that store a magnetic-flux quantum.
- Low loss, low dispersion integrated superconducting transmission lines that support “ballistic” data and clock transfer at the clock rate.
- Cryogenic operating temperatures that reduce thermal noise and enable low power operation.
- RSFQ circuit fabrication that can leverage processing technology and computer-aided design (CAD) tools developed for the semiconductor industry.

Additional discussion of RSFQ technology basics can be found in Appendix D.

1.4 SUMMARY OF PANEL’S EFFORTS

The ITRS has identified RSFQ logic as the lowest risk of the emerging logic technologies to supplement CMOS for high-end computing. The panel’s report:

- Provides a detailed assessment of the status of RSFQ circuit and key supporting technologies.
- Presents a roadmap for maturing RSFQ and critical supporting technologies by 2010.
- Estimates the investment required to achieve the level of maturity required to initiate development of a high-end computer in 2010.

TABLE 1-1. THE PANEL’S CONCLUSIONS:

- Although significant risk items remain to be addressed, RSFQ technology is an excellent candidate for the high-speed processing components of petaflops-scale computers.
- Government investment is necessary because private industry has no compelling financial reason to develop this technology for mainstream commercial applications.
- With aggressive federal investment (estimated between \$372 and \$437 million over five years), by 2010 RSFQ technology can be brought to the point where the design and construction of an operational petaflops-scale system can begin.

1.4.1 RSFQ READY FOR INVESTMENT

In the opinion of the panel, superconducting RSFQ circuit technology is ready for an aggressive, focused investment to meet a 2010 schedule for initiating the development of a petaflops-class computer. This judgment is based on:

- An evaluation of progress made in the last decade.
- Projection of the benefits of an advanced very-large-scale integration (VLSI) process for RSFQ in a manufacturing environment.
- A reasonable roadmap for RSFQ circuit development that is coordinated with manufacturing and packaging technologies.

1.4.2 RSFQ CAN LEVERAGE MICROPROCESSOR TECHNOLOGY

Although RSFQ circuits are still relatively immature, their similarity in function, design, and fabrication to semiconductor circuits permits realistic extrapolations. Most of the tools for design, test, and fabrication are derived directly from the semiconductor industry, although RSFQ technology will still need to modify them somewhat. Significant progress has already been demonstrated on limited budgets by companies such as Northrop Grumman and HYPRES, and in universities such as Stony Brook and the University of California, Berkeley.

Today, individual RSFQ circuits have been demonstrated to operate at speeds in the hundreds of GHz, and system clocks greater than 50GHz seem quite attainable. The devices' extremely low power will enable systems with greatly increased computational capability and power requirements comparable to today's high-end systems.

1.5 ROADMAP CREATED AND GOVERNMENT INVESTMENT NEEDED

Because there is no large commercial demand for superconductive electronics (SCE) technology either currently or expected in the foreseeable future, private industry sees no financial gain in developing it. For this reason, government funding is critical to SC technology's development. Besides its use to NSA, SC will likely have applications for other elements of the government, and once it has been developed, SC may eventually develop commercial applications as well.

While the panel finds that the RSFQ technology is quite promising as a base for future HEC systems, significant developmental effort will be needed to bring it to the state of maturity where it is ready for design and construction of operational systems. The panel believes this technology is at a state of readiness appropriate for a major investment to carry out this development.

To scope out the nature and magnitude of development effort needed, this report presents a detailed technology roadmap, defining those RSFQ technology tools and components that must be developed for use in HEC in the 2010 time frame. The investment required is estimated between \$372 and \$437 million over five years.

The end point of this roadmap defines demonstrations that will validate the technology as ready to be designed into operational systems, including an RSFQ:

- Processor of approximately 1-million gate complexity, on a single multi-chip module (MCM), operating at a 50 GHz clock rate, with inter-chip data communications at the clock frequency.
- Cell library and a complete suite of computer-aided design (CAD) tools.
- Chip manufacturing facility with an established stable process operating at high yields.




Figure 1-2. Roadmap for RSFQ technology tools and components.

1.5.1 FUNDING

Three Funding Scenarios

The development of RSFQ will vary depending on the level of available federal funding. Three possibilities were considered: aggressive government funding, moderate government funding, and a scenario devoid of any additional government investment.

TABLE 1-2. THREE FUNDING SCENARIOS	
Level of Funding	Expected Results
Aggressive Government Funding.	The roadmaps presented assume aggressive government funding.
Moderate Government Funding.	This scenario presumes a sharing of the costs of development between government and industry, which the panel does not believe is a realistic expectation. With reduced government funding and without industrial contribution, the panel sees the technologies maturing much more slowly.
No Additional Investment.	This scenario leaves the entire development to industry. In this case, it is unlikely that the core SCE technology would mature sufficiently to build a machine in the foreseeable future. The panel would expect continued low-level investment in system architecture and programming environments, but no focused effort.

1.6 TECHNICAL ISSUES

1.6.1 SUPERCOMPUTING SYSTEM CONSTRAINTS ON THE USE OF RSFQ

While the goal of this study is not to develop a supercomputer system architecture, architectural constraints and opportunities are important for understanding how superconducting electronics technology would be deployed in a large system and what functionality would be demanded of the superconducting electronics. The earlier Hybrid Technology Multi-Threaded (HTMT) project provided a basis for understanding possible architectures using RSFQ electronics. Appendix F: System Architectures expands on the discussion below. (The full text of this appendix can be found on the CD accompanying this report.)

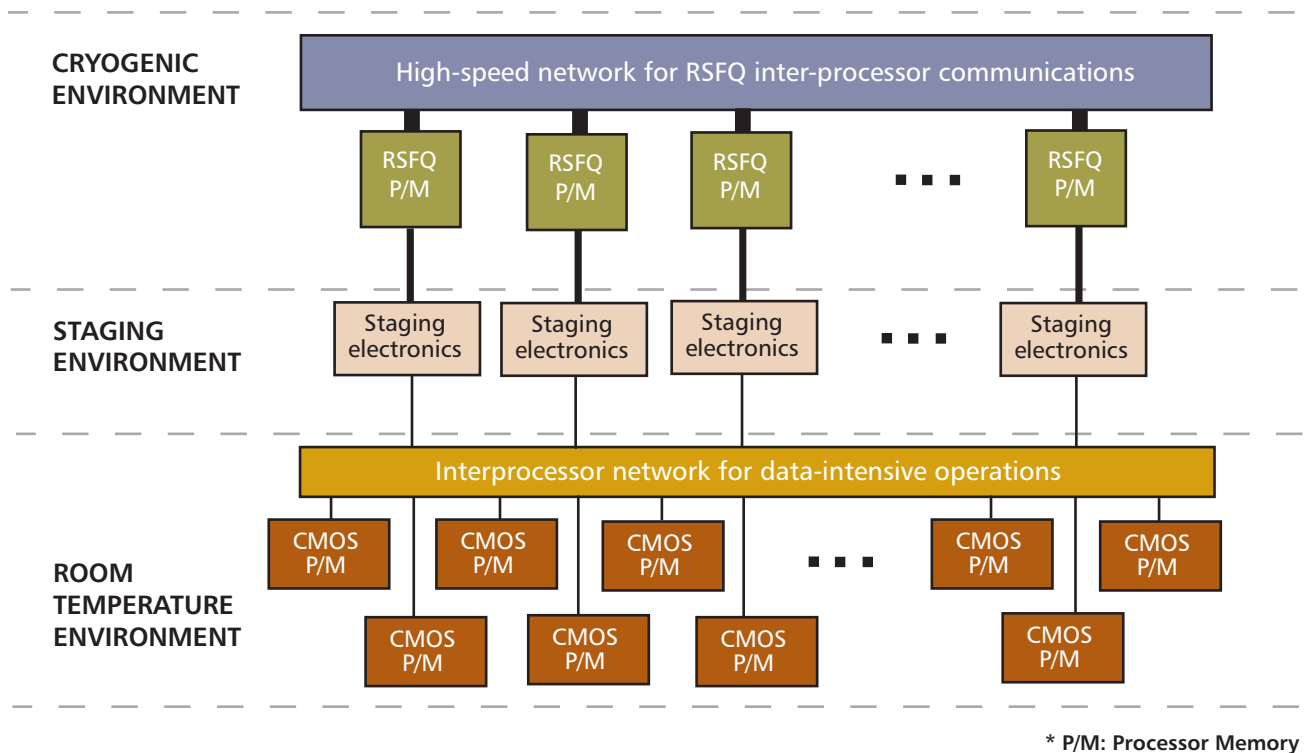


Figure 1-3. Conceptual architecture.

Structurally, a high-end computer will have a cryogenic core to house the RSFQ processors, including a communications fabric linking the RSFQ processing elements; surrounding this will be “staging” electronics connecting the core to a high-bandwidth, low-latency (optical) network; the network will connect to room temperature bulk memory and storage devices. The rationale is that:

- The superconducting processors provide very-high-speed computation, but memory capacity at 4 K is presently limited by low-density. This means that data must be swapped in and out, requiring communications between superconducting processing elements and the outside world to be as fast as possible.
- The staging electronics will orchestrate the movement of data and instructions to and from the RSFQ processors, minimizing stale data in the critical high-speed memory resource. Communication to main memory will need a high-bandwidth network to support large-scale data movements and to provide a way to “gather” dispersed data from main memory into the staging electronics. The key system constraint is the availability and use of cryogenic memory.

1.6.2 FAST, LOW LATENCY MEMORY

Large and fast systems will require fast memory that can be accessed in a small number of machine cycles. As the cycle time decreases, this requirement becomes very demanding, limited by the speed of light as well as the memory speed itself, and presently there is no memory technology that can provide the combination of access time and size that will be needed. The panel has identified three promising approaches to placing significant amounts of fast, low latency RAM at 4 K next to the processor:

- Hybrid JJ-CMOS RAM.
- SFQ RAM.
- Monolithic RSFQ-MRAM.

In addition, CMOS DRAM and MRAM could be located at an intermediate cryogenic temperature (40 - 77 K) to reduce latency. The system architect will then have several options in designing the memory hierarchy.

1.6.3 HIGH SPEED INPUT/OUTPUT

Communicating high-bandwidth data up from the cryogenic environment to room temperature is also a challenge because present drive circuits consume more power than can be tolerated at the low-temperature stage. One approach is to communicate electrically up to an intermediate temperature stage and then optically up to room temperature.

1.6.4 CAD TOOLS

CAD tools and circuit building blocks must be in place so that designs can be done by a competent digital designer who is not a superconductivity expert. These can build on the tools for CMOS circuitry.

1.6.5 REFRIGERATION

Refrigeration is not considered to be a problem. Existing coolers currently used for applications such as superconducting magnets and nuclear accelerators will meet the need.

1.7 STATE OF THE INDUSTRY

Today, expertise in digital superconducting technology resides in only a handful of companies and institutes:

- The Japanese ISTE (International Superconductivity Technology Center) SRL (Superconducting Research Laboratory) is a joint government/industry center that probably has the most advanced work in digital RSFQ anywhere in the world today.
- HYPRES, a small company in Elmsford, NY, is focused entirely on superconducting digital electronics. Its current market is primarily DoD Radio Frequency (RF) applications and related commercial communication systems. It has operated the only full-service commercial foundry in the U.S. since 1983.
- Northrop Grumman (Space Technology) had the most advanced foundry and design capability until it was suspended last year. That division still has a strong cadre of experts in the field, and the company has a continuing research effort in Baltimore, Maryland.

- Chalmers University in Sweden is developing RSFQ technology to reduce error rates in CDMA cell phone base stations.
- Academic research continues at the Jet Propulsion Laboratory and the Universities of California, Berkeley and Stony Brook. Among government agencies, NIST (Boulder), NSA, and ONR have resident expertise.

Some additional companies are working in analog or power applications of superconductivity, but they are not involved in digital devices.

1.8 CONTENTS OF STUDY

The panel organized its deliberations, findings, and conclusions into five topical areas, with each chapter going into detail on one of the following:

- Architectural considerations.
- Processors and memory.
- Superconductor chip manufacturing.
- Interconnect and input/output.
- System integration.

(The attached CD contains the full text of the entire report and all appendices.)

1.9 CHAPTER SUMMARIES

Architectural Considerations

No radical execution paradigm shift is required for superconductor processors, but several architectural and design challenges need to be addressed.

By 2010 architectural solutions for 50-100 GHz superconductor processors with local memory should be available.

Total investment over five-year period: \$20 million.

RSFQ logic at high clock rates introduces unique architectural challenges. Specifically:

- The on-chip, gate-to-gate communication delays in 50 GHz microprocessors will limit the chip area reachable in one cycle to less than 2 mm.
- One can no longer employ clock or data buses.
- Local clocks will have to be resynchronized both on-chip and between chips.

- Superconductor processors will use a partitioned microarchitecture, where processing occurs in close proximity to data.
- In order to achieve high sustained performance, aggressive architectural techniques will be used to deal with memory access latency.

Processors and Memory

The capability in 2010 should be >1 million JJs per cm^2 , implying $>100,000$ gates per cm^2 with clock speed >50 GHz.

The technology should be brought to the point where an ASIC logic designer will be able to design RSFQ chips without being expert in superconductivity.

An advanced 90 nm VLSI process after 2010 should achieve ~ 250 million JJ/ cm^2 and circuit speeds ~ 250 GHz.

Three attractive memory candidates are at different stages of maturity:

- Hybrid JJ-CMOS memory.
 - Single-flux-quantum superconducting memory.
 - Superconducting-magnetoresistive RAM (MRAM).
-

A complete suite of CAD tools can be developed based primarily on corresponding tools for semiconductors.

Total investment over five-year period: \$119 million.

Processors

The complexity of RSFQ logic chips developed to date has been constrained by production facility limitations and inadequate design tools for VLSI. Although all demonstrated digital chips have been fabricated in an R&D environment with processing tools more than a decade older than CMOS, impressive progress has been made. The panel concludes that, given the design and fabrication tools available to silicon technology, RSFQ circuit technology can be made sufficiently mature by 2010 to support development of high-end computers.

Ready for Investment

The panel concluded that with the availability of a stable VLSI chip production capability, RSFQ processor technology will be ready for a major investment to acquire a mature technology that can be used to produce petaflops-class computers starting in 2010. ("Mature technology" means that a competent ASIC digital designer with no background in superconductor electronics would be able to design high-performance processor units.) This judgment is based on an evaluation of progress made in the last decade and expected benefits of an advanced VLSI process in a manufacturing environment. Although RSFQ circuits are today relatively immature, their similarity in function, design, and fabrication to semiconductor circuits permits realistic extrapolations.

Random Access Memory Options

Random access memory (RAM) has been considered the Achilles heel of superconductor logic. The panel identified three attractive options, in decreasing order of maturity:

- JJ-CMOS RAM.
- SFQ RAM.
- RSFQ-MRAM.

To reduce risk, the panel concluded that development should commence for all three at an appropriate level of investment, with continued funding depending on progress.

Each memory approach offers a different performance level that can be a useful complement to RSFQ processors for high-end computing. The roadmap sets forth a baseline plan to rapidly mature hybrid JJ-CMOS RAM and continue development of the other approaches as progress warrants.

Roadmap

The panel defined a roadmap that will demonstrate a 1-million gate RSFQ processor operating at 50 GHz with supporting cryogenic RAM on a single multi-chip module as a milestone validating the potential for application to petaflops-scale computing.

Superconductive Chip Manufacture

By 2010 production capability for high-density chips should be achievable by application of manufacturing technologies and methods similar to those used in the semiconductor industry.

Yield and manufacturability of known good die should be established and costs should be understood.

The 2010 capability can be used to produce chips with speeds of 50 GHz or higher and densities of 1-3 million JJs per cm^2 .

Beyond the 2010 timeframe, if development continues, a production capability for chips with speeds of 250 GHz and densities comparable with today's CMOS is achievable.

Total investment over five-year period: \$125 million.

Niobium-based Fabrication

Significant activity in the area of digital superconductive electronics has long existed in the United States, Europe, and Japan. Over the past 15 years, niobium (Nb)-based integrated circuit fabrication has achieved a high level of maturity. An advanced process has one JJ layer, four superconducting metal layers, three or four dielectric layers, one or more resistor layers, and a minimum feature size of $\sim 1 \text{ }\mu\text{m}$.

Technical Details

Today's best superconductive integrated circuit processes are capable of producing digital logic IC chips with 10^5 JJ/cm². On-chip clock speeds of 60 GHz for complex digital logic and 770 GHz for a static divider (toggle flip-flop) have been demonstrated. Large digital IC chips, with JJ counts exceeding 60,000 have been fabricated. IC chip yield is limited by defect density rather than by parameter spreads. At present, integration levels are limited by wiring and interconnect density rather than junction size, making the addition of more wiring layers key to the future development of this technology.

Panel's Approach

The panel assessed the status of IC chip manufacturing for superconductive RSFQ electronics at the end of calendar year 2004, projected the capability that could be achieved in the 2010 time-frame, and estimated the investment required for the development of RSFQ high-end computers within approximately five years.

Costs

Manufacturing RSFQ IC chips of the required complexity and in the required volumes for petaflops-scale computing will require both recurring costs associated with operation of the fabrication facility, and nonrecurring costs mainly associated with the procurement cost of the fabrication tools and one-time facilities upgrades.

Roadmap Criteria

The roadmap to an SCE IC chip manufacturing capability was constructed to meet the following criteria:

- Earliest possible availability of IC chips for micro-architecture, CAD, and circuit design development efforts. These IC chips must be fabricated in a process sufficiently advanced to have reliable legacy to the final manufacturing process.
- Firm demonstration of yield and manufacturing technology that can support the volume and cost targets for delivery of known good die for all superconductive IC chip types comprising a petascale system.
- Support for delivery of ancillary superconductive thin film technologies such as flip-chip, single-chip, and multi-chip carriers and MCM and board-level packaging for technology demonstrations.
- Availability of foundry services to the superconductive R&D community and ultimately for other commercial applications in telecommunications, instrumentations, and other applications.

Interconnects and System Input/Output

Packaging and chip-to-chip interconnect technology should be reasonably in hand.

Wideband data communication from low to room temperature is a challenge that must be addressed.

Total investment over five-year period: \$92 million.

Essential supporting technologies for packaging, system integration, and wideband communications—particularly wideband data output from the cryogenic to ambient environment— were assessed.

Compact Package Feasible

Thousands to tens of thousands of SC processors and large amounts of memory will require significant levels of cryogenic packaging. A compact system package is needed to support low latency requirements and to effectively use available cooling techniques. The panel's conclusions for supporting technologies were:

- Input/Output (I/O) circuits, cabling, and data communications from RSFQ to room temperature electronics above 10 GHz, reliable multiple-temperature interfaces and the design for larger applications have not been completely investigated.
- Output interfaces are one of the most difficult challenges for superconductive electronics.
- A focused program to provide the capability to output 50 Gbps from cold to warm electronics must overcome technical challenges from the power dissipation of the interface devices at the cryogenic end.

Packaging

The panel noted that:

- The design of packaging technologies (e.g., boards, MCMs, 3-D packaging) and interconnects (e.g., cables, connectors) for SCE chips is technically feasible and fairly well understood.
- A foundry for MCMs and boards using superconducting wiring is a major need.
- The technology for the refrigeration plant needed to cool large systems—along with the associated infrastructure—is in place today.
- Tools and techniques for testing a large superconducting digital system have not been fully addressed yet.

Optoelectronic Components at Low Temperature

An issue which must be thoroughly explored is how well room temperature optical components function in a cryogenic environment, or whether specially designed components will be needed.

Low Power a Two-edged Sword

The low power of RSFQ presents a challenge for data output. There is not enough signal power in an SFQ data bit to drive a signal directly to conventional semiconductor electronics; interface circuits are required to convert the SFQ voltage pulse into a signal of sufficient power. While Josephson output circuits have been demonstrated at data rates up to 10 Gbps, and there is a reasonable path forward to 50 Gbps output interfaces in an advanced foundry process, a focused program is required to provide this capability.

Interconnection Network

The interconnection network at the core of a supercomputer is a high-bandwidth, low-latency switching fabric with thousands or even tens of thousands of ports to accommodate processors, caches, memory elements, and storage devices. The Bedard crossbar switch architecture, with low fanout requirements and replication of simple cells, is a good candidate for this function.

Optical Switching Looks Promising

The challenges imposed by tens of Pbps between the cold and room temperature sections of a petaflops-scale superconducting supercomputer require the development of novel architectures specifically designed to best suit optical packet switching, which has the potential to address the shortcomings of electronic switching, especially in the long term.

System Integration

The design of secondary packaging technologies and interconnects for SCE chips is technically feasible and fairly well understood.

The lack of a superconducting packaging foundry with matching design and fabrication capabilities is a major issue.

Enclosures, powering, and refrigeration are generally understood, but scale-up issues must be addressed.

System testing issues must be addressed.

Total investment over five-year period: \$81 million.

Background

System integration is a critical but historically neglected part of the overall system design. It is usually undertaken only at later stages of the design. System integration and packaging of superconductive electronics (SCE) circuits offer several challenges due to the extremely high clock rates (50-100 GHz) and operation at extremely cold temperatures (4 - 77 K).

Larger Scale Design Untested

The design of enclosures and shielding for SCE systems is technically feasible and fairly well understood. However, these techniques have never been tested for larger applications, such as a petaflops-scale supercomputer, where dimensions are in the order of several meters.

Multiple Packaging Levels and Temperature Stages

Packaging within the cryogenic enclosure (cryostat) requires an ascending hierarchy of RSFQ chips, MCMs containing and connecting the chips, and boards connecting and providing structure for the MCMs. In addition, MCMs and boards will also be needed at an intermediate (40-77 K) for semiconductor electronics for data I/O and possibly for memory.

Thermal Loads

Superconducting processor chips are expected to dissipate very little power, but the cryocooler heat load for a petaflops system from heat conduction through I/O and power lines between the cryostat and room temperature will be very significant. Reducing this heat load by use of narrower or lower conductivity lines is not practical because the large current for powering RSFQ circuits requires low DC resistances, which translates to thick metal lines. High-bandwidth signal I/O requires low-loss, high-density cabling, which also translates to high conductivity signal lines or large cross-section signal lines. Therefore, each I/O design must be customized to find the right balance between thermal and electrical properties.

The major issue to be addressed will be assuring a source of low-cost packaging at affordable cost. Several approaches should be evaluated:

- Develop a superconducting MCM production capability.*
 - Find a vendor willing to customize its advanced MCM packaging process to include superconducting wire layers.*
 - Procure MCMs with advanced normal metal layers for the bulk of the MCM, then develop an internal process for adding superconducting wiring.*
-

MCMs

The design of MCMs for SCE chips is technically feasible and fairly well understood. However, the design for higher speeds and interface issues needs further development. The panel expects that MCMs for processor elements of a petaflops-scale system will be much more complex and require more layers of controlled impedance wiring than those that have been built today, with stringent cross-talk and ground-bounce requirements. Although some of the MCM interconnection can be accomplished with copper or other normal metal layers, some of the layers will have to be superconducting for the low voltage RSFQ signals. Kyocera has produced limited numbers of such MCMs for a crossbar switch prototype. These MCMs provide an example and base upon which to develop a suitable volume production capability for MCMs.

Affordable Packaging a Concern

It is expected that the technology will be available for such packaging, but the major issue to be addressed will be assuring a source of low-cost packaging at affordable cost. Several approaches should be evaluated:

- Develop a superconducting MCM production capability, similar to the chip production capability (perhaps even sited with the chip facility to share facility and some staff costs). Although this is the most expensive approach, it provides the most assured access.
- Find a vendor willing to customize its advanced MCM packaging process to include superconducting wire layers and procure packages from the vendor. Because of the relatively low volumes in the production phase, the RSFQ development effort would have to provide most, if not all, of the Non-Recurring Engineering (NRE) costs associated with this packaging. Smaller vendors would be more likely to support this than large vendors.
- Procure MCMs with advanced normal metal layers for the bulk of the MCM, then develop an internal process for adding superconducting wiring. This is less expensive than approach 1, but it depends on the vendor's process, which may change without notice.

3D Packaging an Alternative

An alternative to planar packaging on MCMs and boards is 3-D packaging. Conventional electronic circuits are designed and fabricated using a planar, monolithic approach in mind with only one major active device layer. More compact packaging technologies can bring active devices closer to each other allowing short Time-of-Flight (TOF), a critical parameter for systems with higher clock speeds. In systems with superconducting components, 3-D packaging enables higher active component density, smaller vacuum enclosures, and shorter distances between different sections of the system. As an example, 3-D packaging will allow packing terabytes to petabytes of secondary memory in a few cubic feet (as opposed to several hundred cubic feet) and much closer to the processor.

Power and I/O Cables Needed

SCE circuits for supercomputing applications are based on DC-powered RSFQ circuits. Due to the low voltage (mV level), the total current to be supplied is in the range of few Amperes for small-scale systems and can be easily increased to kilo-Amperes for large-scale systems. Serial distribution of DC current to small blocks of logic has been demonstrated, and this will need to be accomplished on a larger scale in order to produce a system with thousands of chips. However, the panel can expect that the overhead of current-supply reduction techniques on-chip will drive the demand for current supply into the cryostat as high as can be reasonably supported by cabling. Additionally, high serial data rate in and out of the cryostat is expected for a petaflops-scale system. This necessitates RF cabling that can support high line count to service thousands of processors while maintaining the high signal integrity and low losses required for low bit error rate.

Refrigeration

The technology for the refrigeration plant needed to cool large systems is understood. Small space and commercial cryocoolers are available, but engineering changes are needed to enlarge them for larger-scale systems. One key issue is the availability of manufacturers. Development funding may be needed for U.S. companies to insure that reliable American coolers will be available in the future. Development toward a 10 W or larger 4 K cooler would be desirable to enable a supercomputer with a modular cryogenic unit.

System Testing Required

A petaflops-scale superconducting supercomputer is a very complex system and offers major challenges from a system integrity viewpoint. A hierarchical and modular testing approach is needed. The use of hybrid technologies—including superconducting components, optical components and conventional electronic components and system interfaces with different physical, electrical and mechanical properties—further complicates the system testing. This area requires substantial development and funding to insure a fully functional petaflops-scale system.

02

No radical execution paradigm shift is required for superconductor processors, but several architectural and design challenges need to be addressed.

By 2010 architectural solutions for 50-100 GHz superconductor processors with local memory should be available.

Total investment over five-year period: \$20 million.

ARCHITECTURAL CONSIDERATIONS FOR SUPERCONDUCTOR RSFQ MICROPROCESSORS

2.1 SUPERCONDUCTOR MICROPROCESSORS – OPPORTUNITIES, CHALLENGES, AND PROJECTIONS

Although this report focuses on the state of readiness of Rapid Single Flux Quantum (RSFQ) circuit technology, this chapter will first address the requirements that this technology must satisfy to fill the government's needs for high-end computing (HEC).

Any change in circuit technology always calls for a reexamination of the processor architectures that work well with it. For RSFQ, these adjustments are necessitated by the combination of its high-speed and speed-of-light limitations on signal propagation between logic elements.

Superconductor processors based on RSFQ logic can potentially reach and exceed operating frequencies of 100 GHz, while keeping power consumption within acceptable limits. These features provide an opportunity to build very compact, multi-petaflops systems with 100 GHz 64/128-bit single-chip microprocessors to address the Agency's critical mission needs for HEC.

In order to be able to initiate the design of a superconductor petaflops-scale system in 2010, the following critical superconductor architectural and design challenges need to be addressed:

- Processor microarchitecture.
- Memory.
- Interconnect.

The panel believes it will be possible to find and demonstrate viable solutions for these challenges during the 2005-2010 time frame.

The key characteristics of superconductor processors, such as ultra-high clock frequency and very low power consumption, are due to the following properties:

- Extremely fast (a few-picosecond) switching times of superconductor devices.
- Very low dynamic power consumption.
- Ultra-high-speed, non-dissipative superconducting interconnect capable of transmitting signals at full processor speed.
- Negligible attenuation and no skin effect in on- and off-chip niobium transmission lines.

While no radical execution paradigm shift is required for superconductor processors, several architectural and design challenges need to be addressed in order to exploit these new processing opportunities.

The key challenges at the processor design level are:

■ **Microarchitecture**

- a partitioned organization.
- long pipelines.
- small area reachable in a single cycle.
- mechanisms for memory latency avoidance and tolerance.
- clocking, communication, and synchronization for 50-100 GHz processors and chipsets.

■ **Memory**

- High-speed, low-latency, high-bandwidth, hybrid-technology memory hierarchy.

■ **Interconnect**

- Low-latency on-chip point-to-point interconnect (no shared buses are allowed).
- Low-latency and high-bandwidth for processor-memory switches and system interconnect.

Most of the architectural and design challenges are not peculiar to superconductor circuitry but, rather, stem from the processor circuit speed itself. At the same time, some of the unique characteristics of the RSFQ logic will certainly influence the microarchitecture for superconductor processors.

Pipelines

With their extremely high processing rates, fine-grained superconductor processor pipelines are longer than those in current complementary metal oxide semiconductors (CMOS) processors. The on-chip gate-to-gate communication delays in 50 GHz microprocessors will limit the space reachable in one cycle to ~1-2 mm. A time to read data from local, off-chip memory can be up to 50 cycles, while long-distance memory access and interprocessor synchronization latencies can be easily an order of thousands of processor cycles.

Latency Problem

The sheer size of the latency problem at each design level requires very aggressive latency avoidance and tolerance mechanisms. Superconductor processors need a microarchitecture in which most processing occurs in close proximity to data. Latency tolerance must be used to mitigate costs of unavoidable, multi-cycle memory access or other on-/off-chip communication latencies. Some latency tolerance techniques (e.g., multithreading and vector processing) that are successfully used in current processors can also work for superconductor processors. Other potential aggressive architectural options may focus on computation (threads) migrating towards data in order to decrease memory access latency.

Bandwidth Issues

Petaflops-level computing requires very-high-bandwidth memory systems and processor-memory interconnect switches. In order to reach the required capacity, latency, and bandwidth characteristics, the memory subsystems for superconductor processors will likely be built with both superconductor and other technologies (e.g., hybrid SFQ-CMOS memory). Multi-technology superconductor petaflops systems will need high-speed, high-bandwidth (electrical and optical) interfaces between sections operating at different temperatures.

Successful resolution of these design issues and the continuing development of the superconductor technology will allow us to produce a full-fledged 100 GHz 64/128-bit, million-gate microprocessor for HEC on a single chip.

Government Funding Necessary

Based on the current status of superconductor technology and circuit design, the panel believes that only a government-funded project can address these critical processor design challenges between now and 2010.

Major Processor Goals

The project will have two major goals for processor design:

- Find viable microarchitectural solutions suitable for 50-100 GHz superconductor processors.
- Design, fabricate, and demonstrate a 50 GHz, 32-bit multi-chip, 1-million gate processor with 128 KB local memory integrated on a multi-chip module (MCM).

Table 2-1 summarizes the key opportunities, challenges, and projections for superconductor microprocessors.

TABLE 2-1. OPPORTUNITIES, CHALLENGES, AND PROJECTIONS FOR SUPERCONDUCTOR MICROPROCESSORS		
Superconductor Technology Opportunities	Architectural and Design Challenges	Projections
<ul style="list-style-type: none">– Ultra-high processing rates.– Ultra-wideband interconnect.– Very low power dissipation in the processor.	<p>Microarchitecture:</p> <ul style="list-style-type: none">– 50-100 GHz clocking.– long pipelines.– small area reachable in a single cycle, memory latency. <p>Memory:</p> <ul style="list-style-type: none">– high-speed, low-latency.– high-bandwidth.– hybrid-technology hierarchy. <p>Interconnect:</p> <ul style="list-style-type: none">– low-latency.– high-bandwidth.	<ul style="list-style-type: none">– 100 GHz 64/128-bit single-chip microprocessor for HEC.– Very compact multi-petaflops level computing systems with acceptable power consumption.

2.2 MICROPROCESSORS – CURRENT STATUS OF RSFQ MICROPROCESSOR DESIGN

The issues of RSFQ processor design have been addressed in three projects: the Hybrid Technology Multi-Threaded (HTMT) project, the FLUX projects in the U.S., and the Superconductor Network Device project in Japan (Table 2-2).

TABLE 2-2. SUPERCONDUCTOR RSFQ MICROPROCESSOR DESIGN PROJECTS					
Time Frame	Project	Target Clock	Target CPU Performance (peak)	Architecture	Design Status
1997-1999	SPELL processors for the HTMT petaflops system (US)	50-60 GHz	~250 GFLOPS/CPU (est.)	64-bit RISC with dual-level multithreading (~120 instructions)	Feasibility study with no practical design
2000-2002	8-bit FLUX-1 microprocessor prototype (US)	20 GHz	40 billion 8-bit integer operations per second	Ultrapipelined, multi-ALU, dual-operation synchronous long instruction word with bit-streaming (~ 25 instructions)	Designed, fabricated; operation not demonstrated
2002-2005	8-bit serial CORE1 microprocessor prototypes (Japan)	16-21 GHz local, 1 GHz system	250 million 8-bit integer operations per second	Non-pipelined, one serial 1-bit ALU, two 8-bit registers, very small memory (7 instructions)	Designed, fabricated, and demonstrated
2005-2015 (est.)	Vector processors for a petaflops system (Japan)	100 GHz	100 GFLOPS/CPU (target)	Traditional vector processor architecture	Proposal

2.2.1 SPELL PROCESSORS FOR THE HTMT PETAFLUPS SYSTEM (1997-1999)

The HTMT petaflops computer project was a collaboration of several academic, industrial, and U.S. government labs with the goal of studying the feasibility of a petaflops computer system design based on new technologies, including superconductor RSFQ technology.

Issues

The HTMT RSFQ-related design work focused on the following issues:

- **A multithreaded processor architecture** that could tolerate huge disparities between the projected 50-60 GHz speed of RSFQ processors (called SPELL) and the much slower non-superconductor memories located outside the cryostat; and
- **The projected characteristics of the RSFQ superconductor petaflops** subsystem consisting of ~4,000 SPELL processors with a small amount of superconductor memory (called CRAM) and the superconductor network for inter-processor communication.

Chip Design

The architecture of SPELL processors was designed to support dual-level multithreading with 8-16 multistream units (MSUs), each of which was capable of simultaneous execution of up to four instructions from multiple threads running within each MSU and sharing its set of functional units. However, no processor chip design was done for SPELL processors; their technical characteristics are only estimates based on the best projection of RSFQ circuits available at that time (1997-1999).

2.2.2 20-GHZ, 8-BIT FLUX-1 MICROPROCESSOR (2000-2002)

The 8-bit FLUX-1 microprocessor was the first RSFQ microprocessor designed and fabricated to address architectural and design challenges for 20+ GHz RSFQ processors. The FLUX-1 design was done in the framework of the FLUX project as a collaboration between the SUNY Stony Brook, the former TRW (now Northrop Grumman Space Technology), and the Jet Propulsion Laboratory (NASA).

New Microarchitecture Development

A new communication-aware partitioned microarchitecture was developed for FLUX-1 with the following distinctive features:

- Ultrapipelining to achieve 20 GHz clock rate with only 2-3 Boolean operations per stage.
- Two operations per cycle (40 GOPS peak performance for 8-bit data).
- Short-distance interaction and reduced connectivity between Arithmetic Logic Units (ALUs) and registers.
- Bit-streaming, which allows any operation that is dependent on the result of an operation-in-progress, to start working with the data as soon as its first bit is ready.
- Wave pipelining in the instruction memory.
- Modular design.
- ~25 control, integer arithmetic, and logical operations (no load/store operations).

Chips

The final FLUX-1 chip, called FLUX-1R chip, was fabricated in 2002. It had 63,107 Josephson junctions (JJs) on a 10.35 x 10.65 mm² die with power consumption of ~ 9.5 mW at 4.5 K.

Operation of a one-bit ALU-register block (the most complex FLUX-1R component) was confirmed by testing. No operational FLUX-1R chips were demonstrated by the time the project ended in 2002.

2.2.3 CORE1 BIT-SERIAL MICROPROCESSOR PROTOTYPES (2002-2005)

Several bit-serial microprocessor prototypes with a very simple architecture called CORE1 α have been designed, fabricated, and successfully tested at high speed in the Japanese Superconductor Network Device project. Participants in this project include Nagoya, Yokohama, and Hokkaido Universities, the National Institute of Information and Communications Technology at Kobe, and the International Superconductivity Technology Center (ISTEC) Superconductor Research Lab (SRL) at Tsukuba in Japan. The project is supported by the New Energy and Industrial Technology Development Organization (NEDO) through ISTEC.

Bit-serial CORE1 α Microprocessor Prototype

A CORE1 α microprocessor has two 8-bit data registers and a bit-serial ALU. A few byte shift register memory is used instead of a RAM for instructions and data. The instruction set consists of seven 8-bit instructions.

These microprocessors have extremely simplified, non-pipelined processing and control logic, and use slow (1 GHz) system and fast (16-21 GHz) local clocks. The slow 1 GHz system clock is used to advance an instruction from one execution phase to another. The fast local clock is used for bit-serial data transfer and bit-serial data processing within each instruction execution phase. A CORE1 α 10 chip has $\sim 7,220$ JJs on a 3.4×3.2 mm² die, and a power consumption of 2.3 mW.

Bit-serial CORE1 β Microprocessor Prototype

The next design planned for 2005-2006 will be an “advanced bit-serial” CORE1 β microprocessor (14 instructions, four 8-bit registers, and two cascaded bit-serial ALUs) with a 1-GHz system clock, and a 21-GHz local clock. The CORE1 β 2 microprocessor is expected to have 9,498 JJs, 3.1×4.2 mm² size, and power consumption of 3.0 mW.

2.2.4 PROPOSAL FOR AN RSFQ PETAFLOPS COMPUTER IN JAPAN (EST. 2005-2015)

New Focus on Supercomputing

The Japanese are currently preparing a proposal for the development of an RSFQ petaflops computer. This project is considered to be the next step in the SFQ technology development after the Superconductor Network Device project in 2007. Organizations to be involved in the new project will include the current participants and new members to reflect a new focus on supercomputing. The project is expected to be funded through the Ministry of Education. Table 2-3 shows key target technical parameters of the proposed petaflops system.

New Process

An important element of the new proposal is the development of a new 0.25- μ m, 160 kA/cm² process with nine planarized Nb metal layers by 2010, which would allow fabricating chips with 10-50M JJs/cm² density, and reaching a single-chip processor clock frequency of 100 GHz.


Details

The architecture of the proposed RSFQ petaflops computer is a parallel-vector processor with 2,048 processing nodes interconnected by a network switch to a 200-TB dynamic RAM (DRAM) subsystem at 77 K through 25-GHz input/output (I/O) interfaces. Each node has eight 100-GHz vector CPUs, with each CPU implemented on one chip and having 100 GFLOPS peak performance. Each vector CPU has a 256KB on-chip cache and a 32MB off-chip hybrid SFQ-CMOS memory, the latter being accessible to other processors via an intra-node SFQ switch. As proposed in 2004, the total system would have 16,384 processors with a peak performance of 1.6 petaflops.

TABLE 2-3. JAPANESE PETAFLOPS PROJECT TARGETS (AS PROPOSED AT THE END OF 2004)

Die size	1 cm x 1 cm
Fabrication process	0.25 μm , 160 kA/cm ² Nb process
Processor clock	100 GHz
Processor performance (peak)	100 GFLOPS
On-chip processor cache	256KB
Off-chip memory per processor	32 MB hybrid SFQ-CMOS
Number of system nodes (8 CPUs per node)	2,048
Intra-node processor-memory bandwidth (peak)	800 GB/sec
Total DRAM memory at 77 K	200 TB (100 GB/node)
Total number of processors per system	16,384
System performance (peak)	1.6 petaflops
Power at the 4.2K stage	18 kW
Power of the cryocooler	12 MW

A processor node is composed of eight SFQ processors



SFQ Multi Chip Module

Die size: 10 mm x 10 mm
Module size: 80 mm x 80 mm
Bandwidth between chips: 64-b x 100Gbps

Process	160 kA/cm ² 0.25 μm Nb process
Processor performance	100 Gflops (Clock frequency: 100 GHz)
Cache size	L1: 256kB, L2: 32 MB (per processor)
Processor/memory bandwidth	800 GB/s (per processor)
Processor node performance	800 Gflops
Power at 4.2k	7.3 W

Figure 2-1. An 8-processor node of the proposed Japanese superconductor petaflops system.

Potential Problems

The initial version of the proposal does not adequately address microarchitectural and other important design issues (e.g., memory latency tolerance), which underscores both the lack of expertise in real world processor and system design and the limitations of the bottom-up approach used in the Japanese RSFQ design project. Nevertheless, it is reasonable to believe that significant revisions will be made to the initial version of the proposal to address these as well as the other issues before and after the expected start of the project in 2005.

2.3 MICROPROCESSORS – READINESS

Progress Made

Recent progress in RSFQ processor logic and memory design has created a foundation on which to build the next step towards full-fledged RSFQ processors with a clock frequency of 50 GHz. Table 2-4 shows the readiness of key components and techniques to be used at superconductor processor and system design levels.

TABLE 2-4. READINESS OF SUPERCONDUCTOR PROCESSOR ARCHITECTURE AND DESIGN TECHNIQUES

Design Level	Readiness	Basis	Comments
Data processing modules: clocking, structure, and design	Component validation in relevant environment	HYPRES, NGST-Stony Brook Univ, SRL-ISTEC (Japan)	Integer data processing chips; the logical design of the 32-bit FLUX-2 floating-point multiplier
On-chip processor storage structures: organization, capacity, and bandwidth	Component validation in laboratory environment	NGST-Stony Brook Univ, SRL-ISTEC	Vector registers; small-size register files and small memory in CORE1 and FLUX-1 chips
Single-chip microprocessors: microarchitecture and physical design	Component validation in laboratory environment	SRL-ISTEC, NGST-Stony Brook Univ	8-bit FLUX-1 and CORE1 processor prototype chips; a partitioned, FLUX-1 bit-streaming microarchitecture
Chipsets: microarchitecture, partitioning, physical design, and chip-to-chip communication and synchronization	Experimental demonstration of critical functions in laboratory environment	NGST-Stony Brook Univ	60 Gb/link/sec chip-to-chip data transfer over a MCM, clock resynchronization with FIFOs
Off-chip memory: organization, capacity, latency, and data bandwidth	Analytical and experimental demonstration of critical functions	NEC, HYPRES, UC Berkeley	2-4 kbit SFQ, and 64 kbit hybrid SFQ-CMOS single-bank, non-pipelined memory chips
Latency-tolerant architectures for superconductor processors and systems	Analytical and simulation studies	HTMT petaflops compute feasibility study, the Japanese 2004 proposal for a petaflops system	Architectural techniques studied/proposed: multithreading, prefetching, thread migration and vector processing

2.4 MICROPROCESSORS – ISSUES AND CONCERNS

Key Architecture and Design Issues

The first step in designing high-performance superconductor processors will require finding viable solutions for several key architectural and design issues:

- Clocking mechanisms for 50 GHz processors.
- Long latencies of deep RSFQ pipelines.
- Small chip area reachable in a single cycle.
- On-chip instruction/data storage structures.
- Memory capacity, latency, and bandwidth for hybrid technology memory systems.
- Architectures to avoid and tolerate memory-access latency.

2.4.1 CLOCKING FOR 50 GHZ RSFQ PROCESSORS

It is clear that 50 GHz RSFQ processors will not be able to rely on classical, synchronous clocking with all clock signals having same phase and frequency for on-chip circuits. While purely asynchronous design is not feasible because of the lack of CAD tools, other alternatives include the globally-asynchronous, locally-synchronous design approach (used for the 20-GHz FLUX-1 microprocessor), which assumes that clock signals in different clock domains can be out of phase with each other. The problem of clocking will be even more severe if RSFQ processors are implemented on multiple chips on an MCM. This will require architects and designers to find viable mechanisms of synchronization between on- and off-chip processor and memory components.

2.4.2 LONG PROCESSING PIPELINES

As found during the FLUX project, both integer and floating-point fine-grained RSFQ pipelines are significantly longer than those in CMOS processors. The productivity of pipelines decreases with the increase in the overall processor pipeline depth, because of dependencies within critical loops in the pipeline. Thus, superconductor processor architects need to find fine-grained concurrency mechanisms that exploit multiple forms of parallelism at data, instruction, and thread levels in order to avoid low sustained performance. While some new architectural approaches (such as the bit-streaming architecture developed for FLUX-1) look promising for integer pipelines, efficient solutions need to be found for dealing with latencies of floating-point, instruction fetch and decode, and memory access operations.

2.4.3 ON-CHIP INTERCONNECT, CHIP AREA REACHABLE IN A SINGLE CYCLE, AND REGISTER FILE STRUCTURES

Superconductor microprocessors need to pay much more attention to on-chip interconnect and on-chip storage structures to develop delay-tolerant designs. One of the principal peculiarities of pulse-based RSFQ logic is its point-to-point type communication between gates, which precludes use of shared buses. Also, the extremely high clock rate significantly limits the chip area reachable within a clock cycle time. For 50-GHz RSFQ processors, this means that the distance between the gates involved in collective operations within each pipeline cannot be larger than ~ 1 mm. It should be noted, however, that RSFQ pipelines have significantly finer granularity (i.e., they carry out much less logic within each stage) than CMOS. Thus, the major negative impact of the limited chip area reachable in a single cycle will be not on processing within pipeline stages, but rather on data forwarding between non-local pipeline stages and on data transfer between pipelines and registers.

The microarchitecture of superconductor processors must support a truly localized computation model in order to have processor functional units fed with data from registers located in very close proximity to the units. Such organization makes distributed, multi-bank register structures much more suited for superconductor processors than the monolithic multi-ported register files used in CMOS microprocessors. An example of such a partitioned architecture for integer processing was developed for the 20 GHz FLUX-1 microprocessor, which could be called processing-in-registers.

2.4.4 MEMORY HIERARCHY FOR SUPERCONDUCTOR PROCESSORS AND SYSTEMS

Efficient memory hierarchy design including capacity, latency, and bandwidth for each of its levels is one of the most important architectural issues. Some memory issues have been addressed in simulation and low-speed experiments, but there are only analytical results for multi-technology memory hierarchy for a petaflops system with superconductor processors.

The principal design issues are:

- Technology choices for each level of memory hierarchy.
- Interfacing between different memory levels.
- Latency avoidance/tolerance mechanisms in the processor microarchitecture.

Currently, there are several technologies that should be studied as potential candidates for inclusion into such a memory hierarchy:

- RSFQ.
- JJ-CMOS.
- SFQ-MRAM.
- Semiconductor SRAM and DRAM (at higher temperatures, outside the cryostat).

RSFQ Memory

- Traditionally-designed RSFQ RAM (e.g., the FLUX-1 instruction memory) rely on point-to-point implementation of the bit and word lines with tree structures containing RSFQ asynchronous elements such as splitters and mergers at each node of these trees. Such design has a very negative effect on both density and latency.
- Currently, the fastest type of purely RSFQ memory is a first-in-first-out (FIFO)-type, shift register memory, which has demonstrated high-speed (more than 50 GHz), good density, and low latency at the expense of random access capabilities. This type of memory can be efficiently used to implement vector registers, which makes vector/streaming architectures natural candidates for consideration for superconductor processors.

JJ-CMOS Memory

- The closest to RSFQ memory in terms of speed, but much higher in density, is hybrid JJ-CMOS memory, for which simulation results show latencies of few hundred picoseconds for 64 kbit memory chips. A 50 GHz processor with a reasonably large off-chip hybrid JJ-CMOS memory on the same MCM will need architectural mechanisms capable of tolerating 30-40 cycle memory access latencies.

- Wave pipelining can potentially reduce the cycle time for the hybrid memory by two to three times, if implementation challenges are resolved in the context of practical JJ-CMOS memory design. Preliminary analysis suggests that the memory subsystem with 8-16 multi-bank JJ-CMOS chips can provide memory bandwidth of up to 400 GB/sec for a 50 GHz processor mounted on the same MCM.

SFQ-MRAM

- While it is not yet at the same level of maturity as other memory technologies, SFQ-MRAM could be another option to build the high-capacity, low-power shared memory within the cryostat (the 4.2 K system stage). Among the major advantages of having such high-capacity cryo-memory would be the dramatic decrease in the number of wires and I/O interfaces between the cryostat and the off-cryostat main (semiconductor DRAM) memory subsystem.

2.4.5 MEMORY LATENCY TOLERANT ARCHITECTURES

Aggressive architectural mechanisms need to be used for dealing with very significant memory latencies in superconductor processors and systems.

As discussed above, the issues of latency avoidance and tolerance exist at several levels for superconductor processors:

- On-chip registers
 - ~1-10 cycles.
- Off-chip local memory (mounted on same MCM as the processor)
 - ~20-100 cycles.
- Global distributed shared memory (through the system interconnect all across a system)
 - Many thousand cycles, depending on the memory hierarchy level and the distance.

The potential candidates include techniques that were implemented or proposed earlier, namely:

- Vector/streaming processing.
- Multithreading.
- Software-controlled data prefetching.
- Actively managed memory hierarchies with processors embedded in memory.
- Latency avoidance with computation (threads) migrating towards data.

Scalability a Problem

However, the scalability of many known architectural latency-hiding mechanisms to tolerate latency is limited. For instance, as proven in current CMOS systems, traditional multithreading and vector processing can quite successfully hide memory access latencies of an order of 100 cycles. But these techniques alone will not be able to hide thousands of cycles of global access latencies in systems with 50 GHz superconductor processors. This problem is not peculiar to superconductor processors. With the continuing divergence of processor and memory speeds, designers of semiconductor supercomputer systems will face a similar dilemma in five to ten years.

It is clear that novel processor and system architectures that enhance locality of computation are required to deal with the expected range of latencies in superconductor processors and systems. There is every reason to believe that for 50 GHz processors latency avoidance and tolerance will be key architecture design issues.

2.5 MICROPROCESSORS – CONCLUSIONS AND GOALS

In the opinion of the panel, it will be possible to find and demonstrate viable solutions for these architectural and design challenges in the context of a practical RSFQ processor to be demonstrated by 2010.

In the 2005-2010 roadmap, the processor design goals will be the following:

TABLE 2-5. PROCESSOR DESIGN GOALS 2005-010
1. Find Architectural Solutions for 50-100 GHz Superconductor Processors
<ul style="list-style-type: none">– Identify performance bottlenecks for superconductor processors at the microarchitectural level.– Develop a viable architecture capable of avoiding and/or tolerating expected pipeline and memory latencies.
2. Design a 50 GHz 32-bit Multi-chip Processor with Local Memory with the Following Characteristics:
<ul style="list-style-type: none">– Clock frequency: 50 GHz¹.– Performance (peak): 50 GOPS (integer), 100-200 GFLOPS (single precision floating-point).– Logic complexity: ~ 1M gates.– Local memory capacity (total): 128 KByte (multiple chips/banks).– Local memory bandwidth (peak): 200-400 GB/sec (~2 Bytes per flop), Packaging: one MCM.
3. Demonstrate Functionality and Evaluate Performance of the 50 GHz 32-bit Processor with Local Memory
<ul style="list-style-type: none">– Develop a detailed simulation model and write a set of test and benchmark programs.– Evaluate functionality and performance of the 50 GHz 32-bit processor with local memory on a set of test and benchmark programs.

The architecture and microarchitecture design tasks are a pacing item. They should be initiated as soon as possible in order to develop efficient solutions, and provide first logical specifications of the processor chips for circuit designers as soon as the RSFQ fabrication facility becomes operational.

¹ "As mentioned earlier in this Chapter 2, a 50 GHz superconductor processor may employ multiple clock domains potentially with different local clock frequencies, which could be higher or lower than the 50 GHz clock target for data processing units."

2.6 MICROPROCESSORS – ROADMAP AND MILESTONES

The roadmap below (Figure 2-2) shows the specification and the proposed schedule of the project stages related to the aggressive development of the architecture and design of the 50 GHz RSFQ processor with local memory during the 2005-2010 time frame. The roadmap assumes full government funding.

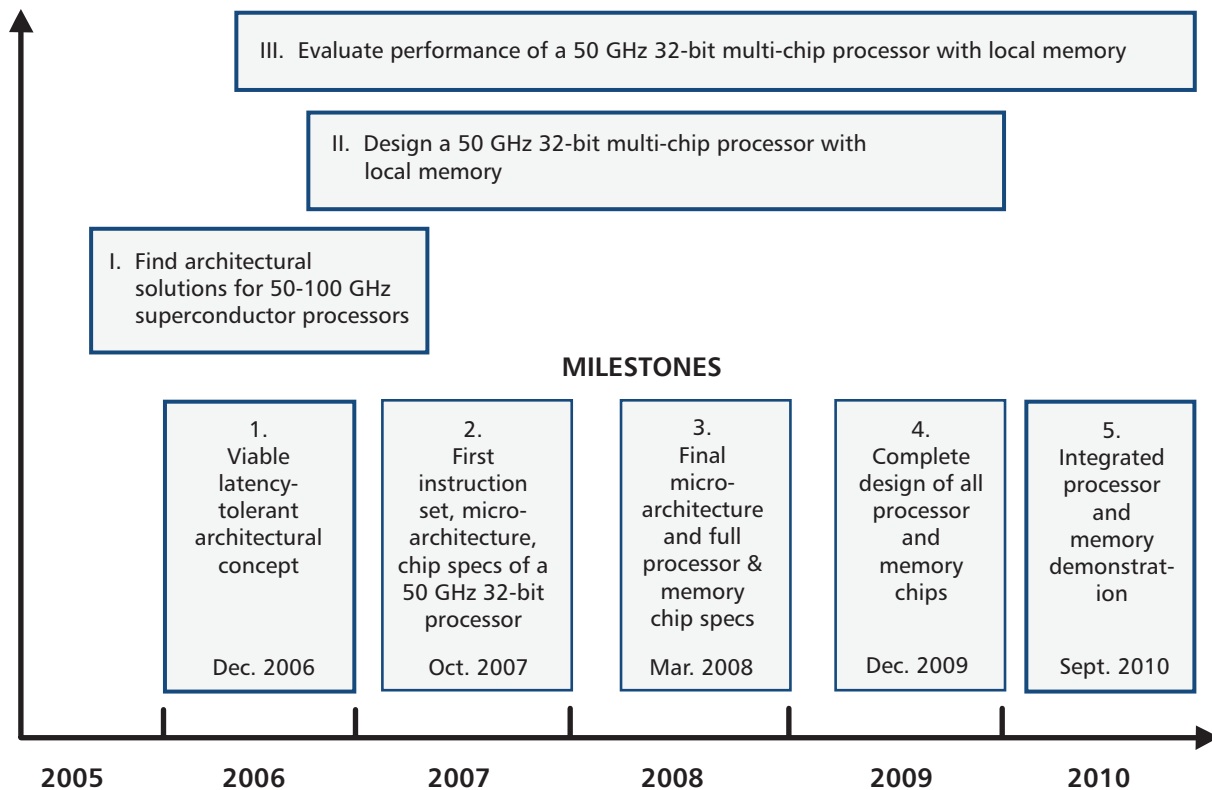


Figure 2-2. 50-GHz multi-chip processor microprocessor design roadmap.

The key milestones for the processor microarchitecture development are:

- **2008:** Microarchitecture of a multi-chip RSFQ processor, with its design simulated and verified for operation at 50 GHz clock frequency.
- **2010:** Integrated functional processor and memory demonstration.

2.7 MICROPROCESSORS – FUNDING

In the table below, “funding” means the funding exclusively for microarchitecture design and evaluation, which does not include physical (circuit-level) chip design, fabrication, and testing.

The total budget for the microarchitecture development and evaluation tasks during the 2005-2010 time frame is approximately \$20 million (M) (Table 2-6). The *Other Expenditures* category includes expenses for purchase of equipment and software (such as PCs, workstations, a multiprocessor simulation cluster, CAD and simulation tools, licenses, etc.), and travel.

TABLE 2-6. FUNDING PROFILE FOR THE MICROARCHITECTURE DESIGN TASKS						
	2005	2006	2007	2008	2009	2010
Viable Architectural concept(s)						
Instruction set, microarchitecture, and processor organization for a 50 GHZ demo						
Chip specifications and gate-level design						
Simulation, performance evaluation, and testing support						
Total man years/year	2	10	10	10	10	10
Est. labor costs/year	\$0.3M	\$2.5M	\$2.5M	\$2.5M	\$2.5M	\$2.5M
Other expenditures	\$0.2M	\$1.5M	\$1.5M	\$1.5M	\$1.5M	\$1.5M
Total for year	\$0.5M	\$4.0M	\$4.0M	\$4.0M	\$4.0M	\$3.5M

03

The capability in 2010 should be >1 million JJs per cm^2 , implying $>100,000$ gates per cm^2 with clock speed >50 GHz.

The technology should be brought to the point where an ASIC logic designer will be able to design RSFQ chips without being an expert in superconductivity.

An advanced 90 nm VLSI process after 2010 should achieve ~ 250 million JJ/ cm^2 and circuit speeds ~ 250 GHz.

Three attractive memory candidates are at different stages of maturity:

- Hybrid JJ-CMOS memory.
 - Single-flux-quantum superconducting memory.
 - Superconducting-MRAM.
-

A complete suite of CAD tools can be developed based primarily on corresponding tools for semiconductors.

Total investment over five-year period: \$119 million.

SUPERCONDUCTIVE RSFQ PROCESSOR AND MEMORY TECHNOLOGY

The panel concluded that superconductive Rapid Single Flux Quantum (RSFQ) processor technology is ready for an aggressive, focused investment to meet a 2010 schedule to initiate the development of a petaflops-scale computer. The panel believes this technology can only be developed with full government funding, since U.S. industry is not moving in this direction and is unlikely to make significant progress toward this goal if left on its own.

“A credible demonstration of RSFQ readiness must include at least three things:

- a foundry that produces advanced chips with high yield,*
- CAD that yields working chips from competent engineers, and*
- output interfaces that send 40 Gbps data from cold to warm electronics.”*

- Dr. Burton J. Smith, Chief Scientist, Cray Corporation

Significant challenges to demonstrating a million-gate, 50-GHz processor with 128 kbyte RAM include:

- Microarchitecture that effectively utilizes the high clock speed.
- RSFQ Computer-Aided Design (CAD) tools and design optimization.
- Adequate very-large-scale integration (VLSI) manufacturing capability.
- Cryogenic testing of high clock rate circuits.
- Multi-layer high-speed superconducting multi-chip module (MCM) substrates.
- Adequate cryogenic RAM.

This chapter:

- Reports the status of superconductive RSFQ processor and cryogenic RAM technology.
- Projects the capability that can be achieved in 2010 and beyond with adequate investment.
- Estimates the investment required to develop RSFQ processors and cryogenic memory for high-end computing (HEC) engines by 2010.

The Hybrid Technology Multi-thread (HTMT) architecture study projected that petaflops performance can be achieved with 4,096 RSFQ processors operating at a clock frequency of 50 GHz and packaged in approximately one cubic meter. Each processor may require 10 million Josephson junctions (JJs). To achieve that goal, RSFQ chip technology must be matured from its present state to ≥ 1 million JJs per square centimeter with clock frequencies ≥ 50 GHz.

The panel categorizes chips required for HEC into four classes that must be compatible in speed, bandwidth, and signal levels:

- Processor.
- Memory.
- Network.
- Wideband input/output (I/O) and communications.

Although these circuits are based on a common device foundation, they are unique and have different levels of maturity.

To provide low latency, a large, fast-access cryogenic RAM is required close to the processors. Since very large bandwidth interconnect is a requirement for HEC, wideband interconnects are also needed at all levels of communications. To achieve these goals, microarchitecture, circuit design, and chip manufacturing need to be improved in synchrony. This chapter is divided into three sections:

- Section 3.1 reviews the status, readiness for major investment, roadmap and associated investment, major issues for RSFQ processors, and future projections.
- Section 3.2 reviews the status, readiness for major investment, roadmap and associated investment, major issues for cryogenic RAM, and future projections.
- Section 3.3 elaborates on CAD tools and design methodology required for successful design of large RSFQ circuits.

3.1 RSFQ PROCESSORS

The panel developed an aggressive, focused roadmap to demonstrate an RSFQ processor with 10 million JJs operating at a clock frequency of 50 GHz by 2010, with 128 Kbytes of fast cryogenic RAM placed on the MCM with the processor.

This section reviews the status, readiness for major investment, roadmap and associated investment, major issues for RSFQ processors, and future projections.




Figure 3.1-1. Roadmap and major milestones for RSFQ processors.

3.1.1 RSFQ PROCESSORS – STATUS

Processor Logic

Only a limited number of organizations are developing RSFQ digital circuitry:

- **HYPRES:** a small company specializing in a wide range of superconductor applications is focused on developing packaged wideband software radio and small mixed signal systems for the government.
- **Northrop Grumman (NG):** is developing both mixed-signal and HEC technology for the government. Until suspending its Nb/NbN fabrication line, it had the most advanced superconductor fab. The firm also teamed with Stony Brook University and JPL to develop the Flux-1 microprocessor chip.

■ **ISTEC-SRL:** a collaboration among government, university, and industry in Japan has mounted the Japanese effort to develop and apply RSFQ technology for high-speed servers and communication routers under a five-year Superconductor Network Device project, funded by the Japanese government. It is currently developing a plan for a petaflops computer project.

■ **Chalmers University:** in Sweden is developing RSFQ technology to reduce the error rates in CDMA cell-phone base-stations.




Figure 3.1-2. Photograph of 1-cm² Flux-1 63K JJ microprocessor chip.

Until recently, only small-scale Single Flux Quantum (SFQ) circuits were being developed for mixed signal applications. The simplest circuits were asynchronous binary ripple counters for A/D converters, but many also included logic sections such as digital filters. Some used multiple low-JJ-count chips flip-chip bonded on a superconducting substrate. Cell libraries were developed by adapting CAD tools developed for the semiconductor industry augmented by tools specialized for superconductors. Complex circuits, incorporating from 7 k to 70 k JJs were designed for operation from 17 to 20 GHz and fabricated in R&D facilities. Many were successfully tested at low speed. Less complex circuits have been successfully tested at speeds in excess of 20 GHz. Table 3.1-1 lists some recently reported circuits. Successful development of RSFQ processors depends critically on the availability of a reliable, industrial-quality VLSI fabrication facility that does not exist today.

TABLE 3.1-1. PUBLISHED SUPERCONDUCTIVE CIRCUITS

Circuits/ Organizations	Function	Complexity	Speed	Cell Library
Flux-1/ NG, Stony Brook U., JPL	8-bit microprocessor prototype 25 30-bit dual-op instructions	63K JJ 10.3x10.6 mm ²	Designed for 20 GHz Not tested	Yes, Incorporates transmission line drivers/receivers
CORE 1_10/ ISTEC-SRL, Nagoya U., Yokohama National U.	8-bit bit-serial prototype microprocessor 7 8-bit instructions 2.3 mW dissipation	7K JJ 3.4x3.2mm ²	21 GHz local clock; 1 GHz system clock Fully functional	Yes. Gates connected by JTL stages and/or striplines
NG	MAC and Prefilter for programmable pass-band ADC	6K - 11K JJ 5x5mm ²	20 GHz	Yes. Gates connected by parameterized JTLs and/or striplines
HYPRES	A/D converter	6 K JJ	19.6 GHz	N/A
HYPRES	Digital receiver	12 K JJ	12 GHz	N/A
NG	FIFO buffer memory	4K bit/ 2.6x2.5mm ²	32 bits tested at 40 GHz	No, Custom memory cells
NSA, NG	X-bar switch	32x32 module	2.5 Gbps	Custom switch cells
NG	SFQ X-bar switch	32x32 module	40 Gbps	Custom switch cells
NG	Chip-to-chip DFQ Driver/Receiver	< 1 k JJ	60 Gbps	N/A
NG	Current recycling	< 1 k JJ	40 Gbps	N/A

All RSFQ circuits to date have been fabricated in an R&D laboratory environment at HYPRES ($J_C = 1 \text{ kA/cm}^2$ and 4.5 kA/cm^2), NEC-SRL ($J_C = 2.5 \text{ kA/cm}^2$), or Northrop Grumman ($J_C = 4$ and 8 kA/cm^2). The number of superconducting wiring layers has been limited to four, including the ground plane; junctions are $\sim 1 \text{ }\mu\text{m}$.

A useful source of the latest results in RSFQ digital technology can be found in the proceedings of the biennial Applied Superconductivity Conference which are published the following year in the IEEE Transactions on Applied Superconductivity. The last Conference was held in September 2004 and the proceedings will be published in the June 2005 IEEE Transactions on Applied Superconductivity.

Embedded Vector Registers, FIFO, and SFQ RAM

In addition to logic, processor chips incorporate local memory. This is frequently in the form of registers and first-in, first-out registers (FIFOs). Small, embedded SFQ RAM is also useful. FIFOs can load and unload data using separate clocks that may be incoherent and are therefore useful as register memory and for resynchronizing data for two circuits running with slightly asynchronous clocks. Northrop Grumman tested a 32-bit RSFQ FIFO at 40 GHz, nearly as fast as the planned 50 GHz processor operation. The FIFO was made in the Northrop Grumman 4 kA/cm² process, where a 4 k FIFO (64 64-bit words) occupies a 5-mm square chip. Initial fabrication improvements, discussed in Chapter 4, will increase J_c to 20 kA/cm², quadrupling the size to 16 kbits (256, 64-bit words). Reducing metal line pitch and adding more superconducting wiring levels will further increase the density. The measured bias margins of the 32-bit FIFO measured at 40 GHz were $\pm 17\%$, compared with $\pm 23\%$ at low speed. Increasing the current density should increase the margins at high speed.

A register file was developed by HYPRES based on a three-port memory cell with one input port and two independent output ports. Two different versions of this cell were demonstrated and patented. Such designs could produce multi-port registers.

Small SFQ memory arrays are used today. Small RAM can more easily approach the clock speed and size required for inclusion on processor chips than large RAM blocks.

Inter-chip Communications

Both active and passive superconducting transmission lines have been used to transmit SFQ data over significant distances on-chip. The active lines are a series of Josephson SFQ repeaters called a Josephson transmission line (JTL). Passive transmission lines (PTL) can propagate SFQ pulses on microstrip or striplines at >60 Gbps. The latter technique is preferred, because it uses less power and produces less jitter. Passive transmission line communication was used extensively on the Flux-1 chip. The velocity in superconducting transmission lines is about 1/3 the free space velocity, 0.1 mm/ps. Time-of-flight (TOF) in JTLs is longer because of the switching delay at every junction.

A more challenging problem is communication of SFQ data between chips on an MCM and, eventually, between PC boards. Communication between chips requires the SFQ pulses to traverse the discontinuities between the chip and the MCM. Transmission lines are matched to the junction impedance, typically 1-10 Ω . On-chip, these lines are several microns wide. The transmission lines on the MCM substrate are much wider, because the dielectric insulation layers are significantly thicker on the MCM. In addition to this geometric discontinuity, the signal and ground solder bumps introduce parasitic inductance and capacitance.

Until 2002, conventional wisdom held that SFQ pulses could not be transmitted through a passive substrate; researchers in Japan demonstrated that SFQ pulses could be transmitted through an “active” substrate. However, Northrop Grumman developed a unique, efficient, high-speed driver-receiver for inter-chip communications through solder-bumps on an MCM and demonstrated low bit-error-rate to 60 GHz. In addition to special communication circuits, it was necessary to tailor the local bump configuration to tune out the parasitics. This required both analog simulations and iterative experimental tests.

The location and distribution of ground bumps and contacts is of equal importance as the signal line. A ground bump introduces inductance not only by the bump geometry, but also by its distance from the signal bump. Ideally, the ground bump would be coaxial with the signal to simulate a coaxial line. Since this is not feasible, Flux-1 used four ground bumps surrounding the signal at minimum separation (Figure 3.1-3). Since many high-speed lines are required, it was necessary to share ground bumps. Simulations showed that cross-talk at the ground bumps was negligible.




Figure 3.1-3. Signal and ground pad geometry used on Flux-1. Pad size and spacing are 100 μm .

3.1.2 RSFQ PROCESSORS – READINESS FOR MAJOR INVESTMENT

Processor

The panel believes that RSFQ processor technology is ready for a major investment to produce a mature technology that can be used to produce petaflops-class computers starting in 2010. “Mature processor technology,” means one that would enable a competent ASIC digital designer, with no background in superconductive electronics, to design high-performance processor units. This judgment is based on an evaluation of progress made in the last decade and projection of manufacturing environments, coupled with a roadmap for RSFQ circuit development coordinated with VLSI manufacturing and packaging technologies. Although large RSFQ circuits are relatively immature today, their similarity in function, design, and fabrication to semiconductor circuits permits realistic extrapolations.

Most of the tools for design, test, and fabrication will continue to be derived from the semiconductor industry, but will need to be adapted for application in RSFQ design. Because of this, investment in this effort will not need to match that needed for semiconductor development. The Flux project at Northrop Grumman, Stony Brook, and JPL illustrates how progress in fabrication and circuit technology can be accelerated in tandem. In the three years from 2000 to 2003, fabrication technology was advanced by two generations, with gate libraries fully developed for each generation, and by early 2004 Northrop Grumman was ready to move to a 20 kA/cm^2 process. This was accomplished on a limited budget. The Flux-1 microprocessor, including scan path logic for testing, was designed and fabricated from a 10-cell library, and inter-chip communication circuits were designed and successfully tested up to 60 Gbps.

Embedded Vector Registers, FIFO, and Small SFQ RAM

Based on published reports, vector registers, FIFOs, and small SFQ RAM arrays are ready for the major investment needed to bring them to the required readiness for HEC. It will be necessary to:

- Enlarge the memory size using an advanced fabrication process (Chapter 4).
- Set up the read/write procedures to meet the needs of the processor.
- Test the FIFO in a logic chip.

Following this, test circuits can be designed and fabricated for testing at 50 GHz. Once problems are resolved, a complete 64 x 64 FIFO register can be designed and fabricated for inclusion in a processor.

Inter-chip Communications

Driver and receiver circuits have already been demonstrated to 60 Gbps per line in the NG fabrication processes. Since speed and bit error rate (BER) will improve as J_C increases, the circuit technology is ready now. The principal task will be to accommodate the increased number of high-data-rate lines on each chip and MCM. Shrinking bump size and minimum spacing, better design tools, and simulation models will be important to minimize reflections and cross-talk at both the signal and ground chip-to-MCM bumps. The crossbar switch program provides a rich experience base in the engineering of ribbon cables. In particular, the trade-off between heat leak and signal loss is well understood.

Table 3.1-2 summarizes the readiness of superconductive circuits today and projects readiness in 2010 based on the roadmap.

TABLE 3.1-2. SUPERCONDUCTIVE CIRCUIT TECHNOLOGY READY FOR MAJOR INVESTMENT AND DEVELOPMENT		
Circuit Type	2004 Readiness	Projected 2010 Readiness
Logic	– Small circuits @ lower speed.	– $>10^6$ JJ/cm ² @ 50 GHz clock.
SFQ RAM	– Experimental 4 kb RAM @ low speed. – Analysis of SFQ ballistic RAM.	– 256 kb SFQ ballistic RAM @ 500 ps access time.
Hybrid RSFQ-CMOS RAM	– Experimental proof-of-concept.	– 256 kb hybrid @ 500 ps.
Monolithic RSFQ MRAM	– RSFQ write/read concept formulated.	– 128 kb hybrid @ 500 ps.
Vector Register, FIFO	– 32 bit FIFO @ 40 GHz.	– 4 kb FIFO @ 50 GHz.
Communication Circuits	– Chip-to-chip @ 60 Gbps.	– 64-bit word-wide chip-to-chip @ 50 GHz.
I/O	– 10 Gbps driver.	– 64-bit word-wide drivers @ 40 Gbps/s.
Switch	– 2.5 Gbps per port circuit switch, scalable.	– 64-bit word wide, scalable @ 50 Gbps per port.

In evaluating the circuit technology and projecting a high probability of success in HEC, the panel considered the availability of VLSI fabrication. As much as anything, success in RSFQ processor development will depend on a dedicated industrial-quality VLSI manufacturing facility that provides a significant step-up in topological design rules and chip yield from current R&D facilities (see Chapter 4).

Cryogenic testing is much more difficult than room temperature testing. Testing at 50 to 100 GHz clock frequencies magnifies the challenge significantly. A high-speed test bench will be required that is customized to test and evaluate all superconductive circuits from the cell libraries through the RSFQ processor integrated with RAM, including memory and network switches. The test bench must provide environmental support (cooling, RF and magnetic shielding), power supplies, commercial and specialized digital test equipment, and fixturing to accommodate various device modules. SFQ diagnostic circuits have been demonstrated and should be employed to meet the requirements of the test plans.

3.1.3 RSFQ PROCESSORS – ROADMAP

The roadmap for RSFQ logic circuits is limited by and dependent on the roadmap for other critical elements of the technology: integrated circuit (IC) fabrication, CAD tools, MCM packaging, and microarchitecture. Early circuit development will be paced by the IC fabrication development.

The approximate sequence of efforts is to:

- Coordinate with the fabrication unit to develop topological design rules.
- Acquire the maximum set of commercially available CAD tools (see Section 3.3).
- Develop and validate a cell library for each set of design rules: first, a gate-cell library and eventually, a macro-cell library.
- Develop a high-speed cryogenic test bench.
- Design, test, and evaluate a complete set of control, register, and processing units, including floating point units, for each set of design rules, guided by the microarchitecture.
- Design, test, and evaluate a 50-GHz embedded memory.
- Evaluate and implement solutions to such issues as margins, jitter, clock distribution, noise, current recycling, etc.
- Integrate 50 GHz processor and memory chips (see Section 3.2) on a single MCM, test, and evaluate the RSFQ processor.

There will be at least two major circuit iterations of increasing speed and complexity paced by the fabrication and architecture roadmaps.

The major milestones are:

TABLE 3.1-3. MAJOR MILESTONES		
Milestone	Time	Cost (\$M)
(start-up efforts underway)	2006	6.9
Cell Library – 1 Completed	2007	9.5
Logic, Control, Register Units – 1 Demonstrated	2007	
Cell Library – 2 Completed	2008	10.6
Logic, Control, Register Units – 2 Demonstrated	2008	
Integrated Processor Unit – 3 Demonstrated	2009	10.6
RSFQ processor with cryo-RAM Demonstrated	2010	8.8
Total Investment		46.4

3.1.4 RSFQ PROCESSORS – INVESTMENT REQUIRED

The panel projects a government investment of \$46.4 million (M) will be required to achieve the processor technology readiness goals to initiate development of a petaflops-class computer in 2010. Table 3.1-4 depicts three investment scenarios. The panel believes that only aggressive government investment will advance the technology toward the requirements of HEC, and that any industrial partnership is unlikely before late in the five-year roadmap.

TABLE 3.1-4. RSFQ PROCESSOR TECHNOLOGY FOR HEC ROADMAP		
Aggressive Funding	Follows roadmap	\$46.4M
Moderate Funding	Negligible technology advance to meet HEC needs over today	Continued low level ONR technology funding focused on Navy problems
No Funding	Negligible technology advance to meet HEC needs over today	No investment

The level of investment is determined by the requirements of the roadmap. Table 3.1-5 summarizes the estimated investment schedule for logic circuit development, including:

- Establishing design rules jointly with the fabrication group.
- Developing cell libraries, chip design and testing.
- Developing a high-speed cryogenic test bench.
- Mounting chips on substrates.

Additionally, a low level of investment in innovative research for future high payoff improvements in density, power, and clock frequency for HEC is essential to the continued availability of RSFQ technology.

TABLE 3.1-5. INVESTMENT PROFILE FOR RSFQ PROCESSOR CIRCUIT TECHNOLOGY (ASSUMING \$250K PER MY)						
YEAR	2006	2007	2008	2009	2010	Total
Labor (\$M)	1.9	3.5	3.6	3.6	3.8	16.4
Other Investment (\$M)	5	6	7	7	5	30
Total Investment (\$M)	6.9	9.5	10.6	10.6	8.8	46.4

3.1.5 RSFQ PROCESSORS – ISSUES AND CONCERNS

Identifying and addressing important issues early will reduce the risk for demonstration of a high-performance RSFQ processor.

TABLE 3.1-6. ISSUES FOR RSFQ PROCESSOR TECHNOLOGY DEVELOPMENT AND RESOLUTION

Issue	Resolution
<ul style="list-style-type: none"> – Margins are impacted by parameter spreads, gate optimization, current distribution cross-talk, signal distribution cross-talk, noise, bias and ground plane currents, moats and flux trapping, and timing jitter. – Shrinking margins have recently been reported for large circuits ($>10^3$ JJ's) for clock rate above 10 GHz. 	<ul style="list-style-type: none"> – Isolate on-chip power bus and data lines from gate inductors by additional metal layers. – Control on-chip ground currents by using differential power supply. – Optimize design for timing.
<ul style="list-style-type: none"> – Low gate density. 	<ul style="list-style-type: none"> – Improve litho to shrink features. – Increase number of superconducting layers so power, data transmission lines, and gates are on separate levels.
<ul style="list-style-type: none"> – Clock skew. – On-chip velocity $\sim 100\mu\text{m/ps}$. – At 50 GHz, distance within clock cycle is 2 mm. 	<ul style="list-style-type: none"> – Increase gate density to reach more gates in clock cycle. – Generate clock on-chip. – Resynchronize local clocks. – Timing tolerant design and microarchitecture.
<ul style="list-style-type: none"> – Thermal and environmental noise increases errors. – All I/O lines are subject to noise pickup and cross-talk. 	<ul style="list-style-type: none"> – Reduce environmental noise by isolating and filtering power and I/O lines.
<ul style="list-style-type: none"> – Timing errors limit clock frequency. – All junctions contribute jitter. – Noise enhances jitter. 	<ul style="list-style-type: none"> – Reduce environmental noise. – Reduce number of JJs in data and clock distribution network. – Improve timing analysis and simulation using VHDL.
<ul style="list-style-type: none"> – Magnetic flux trapped in superconducting films can shift the operating point of devices. – Local magnetic field and large transients sources of trapped flux. 	<ul style="list-style-type: none"> – Improved shielding and filtering. – Develop methodology for design of moats to trap magnetic flux away from circuits.
<ul style="list-style-type: none"> – Dense chips require many amperes of current. 	<ul style="list-style-type: none"> – Supply RF bias-current from RT to ~ 40 K. – RF-DC power converter at ~ 40 K. – Use HTS power leads from 40 to 4 K. – Extensive current reuse on chip and between chips.
<ul style="list-style-type: none"> – No vendor for superconducting MCM. 	<ul style="list-style-type: none"> – Decision between internal or vendor development.

Appendix G: Issues Affecting RSFQ Circuits provides an expanded discussion of these topics. (The full text of this appendix can be found on the CD accompanying this report.)

3.1.6 RSFQ PROCESSORS – PROJECTED FUTURE CAPABILITY

Chip density is a major factor in the ultimate success of RSFQ technology in HEC, and in many other nascent applications. Comparing present 1- μm RSFQ technology with 0.25- μm CMOS that has many more wiring layers, the panel can draw general conclusions about RSFQ density limits. As described further in Chapter 4, the last Northrop Grumman process could be extended to 1 million JJs/cm² in one year with a few readily achievable improvements in the fabrication process.

This does not imply that 1 million JJs/cm² is the limit of RSFQ technology; greater density would be achieved from further advances in fabrication, design, and innovative concepts. For example, self-clocked circuits without bias resistors is one concept that would reduce power and eliminate the separate clock distribution circuitry, thereby reducing clock jitter and increasing gate density. The panel foresees multi-level circuit fabrication multiplying the functionality of chips as another step forward.

3.2 MEMORY

Cryogenic RAM has been the most neglected superconductor technology and therefore needs the most development. The panel identified three attractive candidates that are at different stages of maturity. In order of their states of development, they are:

- Hybrid JJ-CMOS RAM.
- Single-flux-quantum superconductive RAM.
- Superconducting-magnetoresistive RAM (MRAM).

To reduce risk, the panel concluded that development should commence for all three, with periodic evaluation of progress and relative investment.

It is essential that the first-level memory be located very close to the processor to avoid long time-of-flight (TOF) delays that add to the latency. The panel evaluated options for cryo-RAM, including the critical first-level 4 K RAM that could be located on an MCM adjacent to the processor. Table 3.2-1 compares these memory types.

TABLE 3.2-1. COMPARISON OF READINESS AND CHARACTERISTICS OF CANDIDATE CRYOGENIC RAM					
Memory Type	Readiness for Investment	Potential Density	Potential Speed	Potential Power Dissipation	Cost (\$M)
Hybrid JJ-CMOS	High	High	Medium	Medium	14.2
SFQ	Medium	Medium	Medium-high	Low-medium	12.25
MRAM (40 K)	Medium	High	Medium	Medium	22
SC-MRAM (4K)	Low	High	High	Low	
Monolithic RSFQ-MRAM	Very low	High	Very high	Very low	
Total Investment					48.45

Since these memory concepts are so different from each other, they will be discussed separately. Much of the focus in this study was on the 1-Mb memory to be integrated with the processor in the 2010 milestone. However, it is understood that in a memory hierarchy, this will be backed up by a much larger memory located at 40-77 K, for which MRAM and DRAM are candidates. At 77 K, CMOS channel leakage is negligible, so static power dissipation would be very low and retention time effectively infinite, allowing small DRAM-type cells to be used without refreshing. MRAM power dissipation from the ohmic lines may be substantially reduced at these temperatures.

More details on the MRAM memory technology can be found in Appendix H: MRAM Technology for RSFQ High-End Computing. (The full text of this appendix can be found on the CD accompanying this report.)

The status of the various low-latency cryo-RAM approaches appears in Table 3.2-2.

TABLE 3.2-2. STATUS OF LOW-LATENCY CRYO-RAM CANDIDATES

Type/Lab	Access Time	Cycle Time	Power Dissipation	Density	Status
Hybrid JJ-CMOS (UC Berkeley)	500 ps for 64 kb	0.1 - 0.5 ns depending on architecture	12.4 mW read 10.7 mW write (Single cell writing)	64 kb in < 3x3 mm ²	All parts simulated and tested at low speed
RSFQ decoder w/ latching drivers (ISTEC/SRL)	?	0.1 ns design goal	107 mW for 16 kb (Estimate)	16 kb in 2.5 cm ² (Estimate*)	256b project completed (Small margins)
RSFQ decoder w/ latching drivers (NG)	?	2 ns	?	16 kb/cm ² *	Partial testing of 1 kb block
SFQ RAM (HYPRES)	400 ps for 16 kb (Estimate)	100 ps for 16 kb (Estimate)	2 mW for 16 kb (Estimate)	16 kb/cm ² *	Components of 4 kb block tested at low speed
SFQ ballistic RAM (Stony Brook University)	?	?	?	Potentially dense Requires refresh	Memory cell and decoder for 1 kb RAM designed
SFQ ballistic RAM (NG)	?	?	?	Potentially dense Requires refresh	SFQ pulse readout simulated
MRAM (40K)	Comparable to hybrid CMOS	Comparable to hybrid CMOS	< 5mW at 20GHz (Estimate)	Comparable to DRAM (Estimate)	Room temperature MRAM in preproduction; Low temperature data sparse

*Densities of JJ memories are given for the technologies in use at the time of the cited work. Greater densities can be expected when a 20 kA/cm² process is used. The symbol ? signifies insufficient data.

3.2.1 MEMORY – HYBRID JOSEPHSON-CMOS RAM

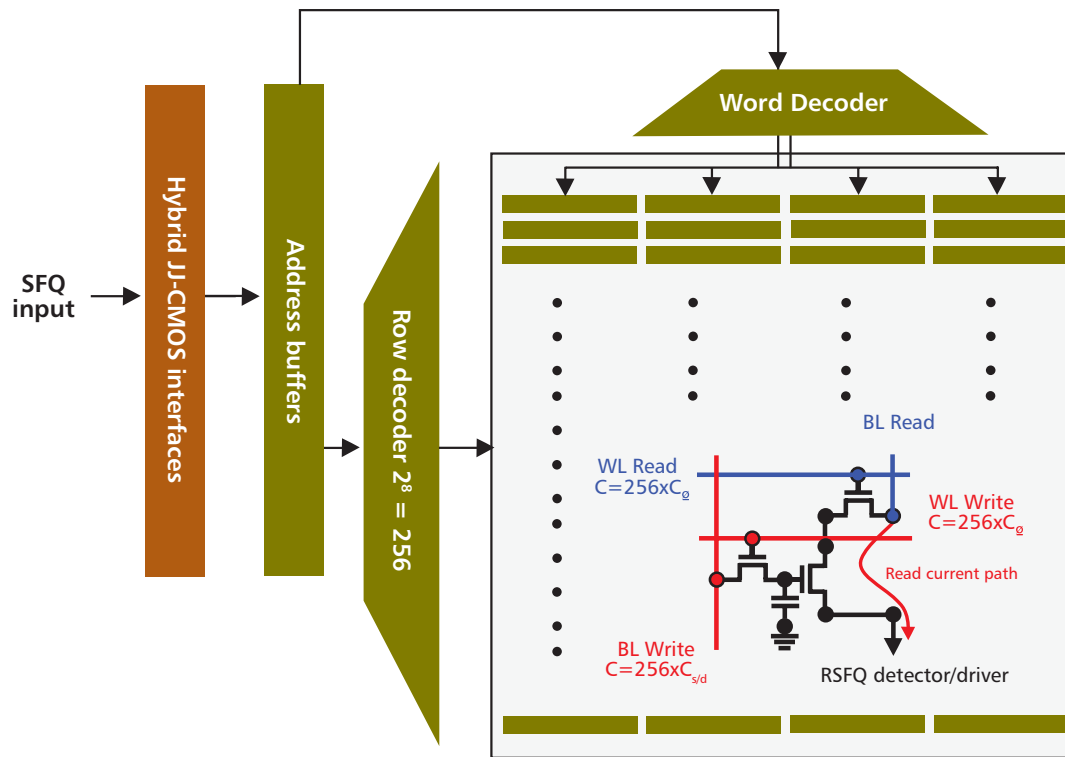


Figure 3.2-1. Hybrid Josephson-CMOS RAM operates at 4 K. The input and output signals are single-flux-quantum pulses.

The hybrid JJ-CMOS RAM uses CMOS arrays at 4 K for data storage, combined with JJ readout to reduce READ time and eliminate the power of the output drivers. In order to access the CMOS, it is necessary to amplify SFQ signals to volt levels. Extensive simulations and experiments have been performed at UC Berkeley on this approach for several years. Furthermore, the core of the memory is fabricated in a CMOS foundry and benefits from a highly developed fabrication process, and the Josephson parts are rather simple. This combines the advantages of the high density achievable with CMOS and the speed and low power of Josephson detection. The entire memory is operated at 4 K, so it can serve as the local cryogenic memory for the processor. A 64-kb CMOS memory array fits in a 2 mm x 2 mm area. As CMOS technology continues to develop, the advances can be incorporated into this hybrid memory. The charge retention time for a three-transistor DRAM-type memory cell at 4 K has been shown to be essentially infinite, so that refreshing is not required. The operation is as though it were an SRAM cell, even though DRAM-type cells are used. Figure 3.2-1 illustrates the overall architecture.

Hybrid JJ-CMOS RAM Status

The experimental part of this work used standard foundry CMOS tested at 4 K. CMOS circuits from four different manufacturers. A BSIM (a simulation tool developed at the University of California, Berkeley) model (industry standard at 300 K) was developed for 4 K operation and gives very good agreement with ring-oscillator measurements, which show a speed-up of ~25% by cooling from 300 to 4 K. It is inferred that the decoder/driver circuits are similarly enhanced upon cooling.

Josephson circuits are used only at the input and output. Very fast, extremely low power, low impedance Josephson detectors are used to sense the bit lines and generate SFQ output pulses. The input requires amplification of the mV SFQ signals to volt-level input to the CMOS driver/decoder. Several circuits have been evaluated; a hybrid combination of Josephson and CMOS components is presently used.

Simulations show that the input interface amplification delay can be <100 ps. Combining the advanced 20 kA/cm^2 JJ process and 90 nm CMOS, the estimate is 75 ps. Parts of the interface circuit have been demonstrated at low speed; high-speed measurements are expected in 2005. The CMOS memory core with SFQ bit-line detectors has been successfully tested at low-speed.

Power dissipation is reduced at 4 K compared with room temperature. CMOS logic dissipates no power except when switching. In addition, the power loss in CMOS at 300 K due to charge leakage is absent at 4 K . Assuming a $0.25 \mu\text{m}$, 64-kb CMOS RAM with a 1 ns cycle time and a JJ 6.5 kA/cm^2 process, READ power dissipation is estimated at 12.4 mW , single bit WRITE at 10.7 mW . Additional power for a 256 kb memory is estimated at $<1 \text{ mW}$ for reading or writing single bits. Scaling to advanced CMOS technologies will reduce power, but accurate estimates are impossible without the BSIM models, which are not yet available from the CMOS foundries.

*Scaling to advanced CMOS technologies will reduce power,
but accurate estimates are not yet possible.*

Assuming $0.25 \mu\text{m}$ CMOS and 6.5 A/cm^2 JJs, access time is estimated to be 500 ps for 64 kb memory and 650 ps for a 256 kb memory. For 90 nm CMOS and 20 kA/cm^2 JJ technologies, the delays scale down by about a factor of two. With these advanced technologies, a 256 kb RAM should fit on a $5 \text{ mm} \times 5 \text{ mm}$ chip.

Hybrid JJ-CMOS RAM Readiness for Investment

Based on the successful work reported to date, hybrid JJ-CMOS RAM is ready for a significant investment to develop low latency 4 K RAM. A 64-kb memory would be first to be evaluated, followed by larger RAM chips at 128 and 256 kb . Advantage can be taken of scaled CMOS and higher current density JJs as they become available in order to improve the access time and reduce power dissipation.




Figure 3.2-2. Wafer with embedded CMOS memory chips for direct wiring of Josephson peripheral circuits. Other packaging schemes also are possible.

The very compact CMOS memory cells will permit at least 1 Mb on a 1-cm^2 chip. The advantage with regard to power dissipation is that amplification of SFQ pulses to clocked volt-level pulses is independent of the size of the memory, so this power is amortized over larger RAM. Connection of RAM to the processor to achieve sufficient bandwidth and access time would have to be explored.

Cryogenic CMOS memory is also a candidate for 40 K application as a second-level RAM. A key advantage is that the CMOS memory cells do not leak at 40 K, which means that compact 3-transistor DRAM-type memory cells can be used as SRAM without the overhead of refreshing that is necessary at room-temperature. It is therefore possible to make very compact, high-capacity memory located at the 40 K stage as a back-up for 4 K RAM. The memory chip design will require close collaboration between the architecture and memory design teams in the first year of the project in order to find the best techniques for processor-memory communication and RAM chip internal organization to meet the read/write latency, access cycle time, and bandwidth requirements of 50 GHz processors.

Hybrid JJ-CMOS RAM Issues and Concerns

The measurements to date on the hybrid Josephson-CMOS memory have all been done at low speed. It is possible that some unexpected issues relating to the frozen-out substrate in the CMOS could cause difficulty in high-speed operation.

Hybrid JJ-CMOS RAM Roadmap

TABLE 3.2-3. HYBRID JJ-CMOS RAM ROADMAP		
Milestone	Year	Cost (\$M)
Test 64-kb bit-slice, optimize design	2006	2.4
Develop embedded-chip packaging	2006	
Fabricate embedded JJ-CMOS chips	2006	
Design and test 64 kb RAM	2007	2.4
Develop 256 kb RAM, measure access- and cycle-time	2007	
Develop new CMOS decoder and input interface	2008	3.5
Fabricate JJ circuits on whole wafer 90 nm CMOS RAM wafers	2008	
Demonstrate 64 kb and 256 kb RAM chips	2008	
Fabricate JJ circuits on redesigned 90 nm CMOS RAM wafers	2009	3.5
Develop processor-memory communication	2009	
Complete the processor-memory package, test and evaluate	2010	2.4
Total Investment		14.2

Hybrid JJ-CMOS RAM Investment

The \$14.2 million (M) total investment required for hybrid JJ-CMOS RAM detailed above is based principally on a hybrid JJ-CMOS team of about 8 persons for five years plus the cost of fabricating 180-nm CMOS chips and 90-nm CMOS wafers. Part of the initial effort could be carried out in a university in parallel with the main industrial effort. This effort would also be able to evaluate CMOS RAM for the second-level memory. The JJ foundry costs are included in the foundry budget.

3.2.2 MEMORY – SINGLE-FLUX-QUANTUM MEMORY

SFQ RAM is a completely superconductive memory, including decoders, drivers, cells, and readout. The ideas are compatible with RSFQ logic and employ that logic for the decoders and detectors.

SFQ RAM Status

The results of five completely superconductive RAM projects are summarized in Table 3.2-4.

TABLE 3.2-4. SUPERCONDUCTIVE RAM PROJECTS			
Organizations	Project	Parameter	Comments
ISTEC-SRL	<ul style="list-style-type: none"> – RSFQ decoder activates latching drivers, which address SFQ memory cells. – VT cells. 	<ul style="list-style-type: none"> – 256 bits. – 3747 JJ. – 1.67 mW@ 10kHz. – 2.05mmx1.87mm. 	<ul style="list-style-type: none"> – Planned 16 kb of 64 blocks on 2.5cm². – 107 mW for 16 kb. – SRL 4.5 kA/cm² process. – Discontinued.
Northrop Grumman	<ul style="list-style-type: none"> – RSFQ input/output. – NEC VT cells. – Latching driver impedance matched to word transmission line for BW. 	<ul style="list-style-type: none"> – Speed > 1 GHz expected, limited by TOF. – 16 kb on 1-cm² chip in 2 kA/cm² process. – Incomplete. 	<ul style="list-style-type: none"> – Decoder tested @ 0.5 GHz. – 1 kb RAM with address. line and sense amplifiers partially tested. – 6 x 6 array tested at low speed, 1-cell @0.5 GHz.
HYPRES	<ul style="list-style-type: none"> – CRAM. – Pipelined DC powered SFQ RAM. 	<ul style="list-style-type: none"> – Long junctions used. – 400ps access, 100ps cycle time estimated. – 16 kb on 1-cm². – Power estimated @ 8 mW for 16 kb. 	<ul style="list-style-type: none"> – 16 kb design used 4 4kb subarrays. – 4 kb block tested at low speed. – Due to block pipeline, time for 64-kb RAM scales as: 600ps access, ~100ps cycle time. – Density increase, cycle time reduced to 30ps, access time ~400ps projected with 20kA/cm². – Discontinued at end of HTMT project.
Stony Brook University	<ul style="list-style-type: none"> – RAM with SFQ access. 	<ul style="list-style-type: none"> – Read pulse travels on active line of JTL/cell stages. 	<ul style="list-style-type: none"> – Developed SFQ cell and decoder for 1kb RAM. – Problem is strong content dependence affected operation.
Northrop Grumman	<ul style="list-style-type: none"> – Ballistic RAM (BRAM). – Cells are 1-JJ, 1-inductor. – SFQ pulses not converted to voltage-state drive levels. 	<ul style="list-style-type: none"> – SFQ pulses propagate through bit lines and are directly detected at end of bit line. – Bit lines are controlled impedance passive transmission lines. – Waveform generated on word line couples ~50mA to each cell in row. 	<ul style="list-style-type: none"> – Simulated SFQ signals passed through series of 64 cells in ~50ps to simulate read. – Decoder delay estimated at 40ps. – DRO cells so refresh must be accommodated.

In summary, several ideas have been studied. In two of these, the memory cells were the NEC vortex transitional (VT) type cells, which were developed for use with voltage-state logic and require voltage-state drivers. One design provided for SFQ input and output pulses on a 256-bit array with potential for high speed. However, scaled to 64 kb, this array would have untenable power dissipation and large size. The other group did not demonstrate a complete system.

The cryo-RAM (CRAM) developed at HYPRES had attractive estimated access and cycle times for 4-kb blocks, but no high-speed data. The estimates for power and access and cycle times scaled to 64-kb are attractive. Circuit density and cycle and access times would be more favorable with increased current density. The ballistic RAM (BRAM) is an interesting concept that will require more research. For all, more advanced fabrication would lead to greater circuit density.

SFQ RAM Readiness for Investment

Although SFQ RAM technology is not as advanced as hybrid CMOS, the panel feels that very attractive concepts exist and should be investigated to support RSFQ processors for HEC. An effort should be made to further evaluate the CRAM and BRAM in order to project their advantages and disadvantages and to quantify power, access time, density, and probability of success. The panel urges that the better choice be pursued at least to the bit-slice stage to further define the key parameters. If successful, the development should be completed.

SFQ RAM Projections

The CRAM design should be scaled to a 20 kA/cm² process and 64-kb RAM to determine the important parameters of speed, power, and chip size. The BRAM concept will require more design evaluation to determine its potential, but either of these ideas could have advantages over the hybrid Josephson-CMOS memory in speed and power.

As noted for the hybrid memory, memory chip design will require close collaboration between the architecture and memory teams in order to find the best processor-memory communication and internal memory chip organization to meet the read/write latency, access cycle time, and bandwidth requirements of 50 GHz processors.

SFQ RAM Issues and Concerns

As in any large RSFQ circuit, margins and flux trapping may be problems. The testing of the CRAM was only done at low speed, so new issues may arise when operated at high speed. In the present DRO BRAM concept, methods for refreshing will have to be explored in order to retain the beneficial aspects of this memory.

SFQ RAM Roadmap

TABLE 3.2-5 SFQ RAM ROADMAP AND INVESTMENT PLAN

Milestone	Year	Cost (\$M)
Design 64-kb CRAM and BRAM and design bit slice for test and evaluation	2006	2.45
Test and evaluate bit slice	2007	2.45
Design, test, and evaluate complete 64 kb memory	2008	2.45
Interface memory with processor	2009	2.45
Integrate memory with processor; Test and evaluate	2010	2.45
Total Investment		12.25

Investment for SFQ Memory

The investment plan (Table 3.2-5) for development of SFQ RAM reflects only manpower costs for design and test. Masks and fabrication costs are in the foundry budgets, the cryogenic high-speed test bench is included in the section 3.1 plan, and RAM-unique test equipment is shared with hybrid JJ-CMOS and MRAM. The total investment for SFQ RAM is \$12.25 million.

3.2.3 MEMORY – MRAM

The panel identified three options for cryogenic MRAM that differ in speed, power, and maturity. RSFQ-MRAM at 4 K could provide the large, low-latency cryogenic RAM required to support RSFQ processors in HEC applications.

TABLE 3.2-6. OPTIONS FOR CRYOGENIC MRAM			
Option	Speed	Power	Maturity
RSFQ-MRAM at 4 K	Very High	Very Low	Very Low
SC-MRAM at 4 K	High	Low	Low
MRAM at 40 K	Moderate	Moderate	Medium

Background

MRAM bits are small ferromagnetic stacks that have two stable magnetic states, retain their value without applied power, and are sensed by coupled resistive elements. MRAM relies on spintronics, where the magnetic polarization of the films affects the multilayers coupled to the ferromagnets. MRAM combines spintronic devices with CMOS read/write circuitry to deliver a combination of attributes not found in any other CMOS memory: speed comparable to SRAM, cell size comparable to or smaller than DRAM, and inherent nonvolatility independent of an operating temperature or device size. Industry is presently developing room temperature MRAM for high performance, nonvolatile applications.

The panel evaluated MRAM that use two different spintronic devices:

- Field switched (FS) tunneling magnetoresistive (TMR) devices.
- Spin momentum transfer (SMT) giant magnetoresistive (GMR) metallic devices.

Both devices rely on spin-polarization to produce a magnetoresistive ratio (MR) that can be distinguished as “0” and “1” with low bit error rate (BER).

Table 3.2-7. COMPARISON OF FIELD SWITCHED TUNNELING MAGNETORESISTIVE AND SPIN MOMENTUM TRANSFER MRAM

Field Switched Tunneling Magnetoresistive MRAM	Spin Momentum Transfer MRAM
<ul style="list-style-type: none"> – Bit is based on tunneling through a thin insulator encased by two ferromagnetic films – Large current-write-pulse produces magnetic field that flips the polarization of one ferromagnetic film – Read by smaller current through the MTJ whose resistance depends on the relative polarization of the two magnetic films – Typical resistances are 10's of $k\Omega$, compatible with CMOS – Commercial pre-production at Freescale 	<ul style="list-style-type: none"> – Bit is low resistance GMR metal – Bit magnetoresistance is changed by spin momentum transfer from write current directly into the film – Read by measuring the resistance of GMR film – Typical GMRs are 1 – 100W, compatible with RSFQ and favored for hard disc drives and magnetic sensors – Opportunity to monolithically integrate fast, low power RSFQ circuits with high speed, high density SMT MRAM operating at 4 K.

MRAM Technology Status

FS-TMR MRAM

- 4Mb FS-TMR MRAM chip, $1.55 \mu m^2$ cell, $0.18 \mu m$ CMOS is in pre-production at Freescale Semiconductor.
- Write currents 1-10 mA, ~25 ns symmetric read-write cycle.
- Circuit architectures can be optimized for lower power, higher speed, or higher density.
- IBM has recently demonstrated a 128-kb, 6 ns access-time MRAM circuit.

4MB MRAM BIT CELL: 1 MTJ & 1 TRANSLATOR




Figure 3.2-3. Diagram of FS-TMR cell (Courtesy of J. Slaughter, Freescale Semiconductor)

In practice, a relatively small increase in MR can significantly improve read-performance and, thus, memory access. As a result, research is focused on higher MR materials. Recent publications have reported that MgO has four-times higher MR than the present AlOx tunnel barrier in TMR cells. Sub-ns memory access is expected, especially at cryogenic temperatures, from continuing improvements in control of thin film micromagnetics. In contrast, write times below 500 ps are unlikely at any temperature in TMR devices, since they are limited by the fundamental Larmor frequency and bit-to-bit variability. Extrapolating performance of TMR FS-MRAM to cryogenic temperatures appears promising.

SMT MRAM

SMT MRAM has only one read/write current path and can yield much higher densities than FS-TMR. A nanomagnetic bit has recently been demonstrated. Cornell University, NIST, and Freescale have demonstrated that pulse current densities of 10^7 - 10^8 A/cm² can rotate the magnetic polarization at 40 K and switching times decrease from 5 ns to ~600 ps at high pulse amplitudes. SMT switching is most effective for single magnetic domains, ideal for bits smaller than ~300 nm. Present SMT structures rely on GMR with low MR, but higher MR materials that will reduce the switching current and time have been identified. A 2- Ω , 90-nm device was simulated to switch at 0.1 mA, equivalent to a voltage bias of 0.2 mV, and a simulated device switched in 50 ps.

SMT memory may be more compatible with RSFQ logic than FS devices, but it is even more embryonic and considerable effort will be required to optimize the materials and cell architectures. It should be emphasized, however, that the opportunity afforded by combining MRAM with superconductive circuits would remain unexplored unless it is funded by government investment for this ultra-high performance computing niche.

MRAM Roadmap

The panel considered a three-step sequence to implement MRAM for cryogenic operation, concluding with monolithic RSFQ-MRAM chips at 4 K.

Table 3.2-8. THREE-STEP PROCESS TO IMPLEMENT MRAM FOR CRYOGENIC OPERATION	
Step	Action
FS-TMR MRAM chips at 40 to 70 K	Reduce the resistance of the wiring and increase the signal-to-noise ratio. As a result, one may be able to implement a very large memory in the cryostat, alleviating some burden on the I/O link between 4 K and ambient.
SC-MRAM at 4 K	Replace all MRAM wiring with superconductors and operate at 4 K, eliminating the dominant ohmic loss in the wiring and further improving signal-to-noise ratio.
RSFQ MRAM at 4 K	Replace all CMOS circuits in SC-MRAM with RSFQ by monolithically integrating RSFQ circuitry with MRAM cells, producing dense, low-latency, low-power RSFQ-MRAM chips that can be placed adjacent to the RSFQ processors.

The roadmap identifies early analytical and experimental trade-offs between GMR and TMR materials and FS versus SMT switching at cryogenic temperatures for optimum memory access. This will help bin material options for the different options. The goal is to have all technology attributes defined and the most promising MRAM cell architectures identified in 2008 and preliminary cell libraries completed in 2009. A 32kB SC-MRAM chip is projected for 2010 and an integrated RSFQ-MRAM chip could be available in 2011-2012. It is anticipated that SMT-cell optimization will continue in the background. Detailed milestones are listed below.

The roadmap and investment plan for SC-MRAM and monolithic RSFQ-MRAM fabrication process development are based on an early partnership between the RSFQ technology development and prototyping team and an MRAM vendor, in order to leverage prior private sector investment in MRAM technology even during the program's R&D phase. The estimated investment does not include building and equipping an RSFQ-MRAM pilot line.

MRAM Major Milestones and Investment Plan

TABLE 3.2-9. MRAM ROADMAP AND INVESTMENT PLAN		
Milestone	Year	Cost (\$M)
Establish collaborations with R&D and commercial organizations	2006	2
Test and verify FS-MRAM performance at 77 K	2006	
Fabricate, test, and evaluate FS and SMT devices	2007	4
Optimize cell architectures: GMR/TMR and FS/SMT	2007	
Develop SC-MRAM process module	2007	
Demonstrate SC-MRAM cells	2007	
Design, fabricate, test, and evaluate 32kB SC-MRAM	2008	5
Define monolithic RSFQ-MRAM fabrication process and test vehicle	2008	
Demonstrate individual RSFQ MRAM cells	2009	5
Define JJ-MRAM-microprocessor test vehicle	2009	
Design, fabricate, test, and evaluate 1Mb SC-MRAM	2009	
Demonstrate functional 32 kB RSFQ MRAM	2010	6
Total Investment		22

MRAM Future Projections

MRAM is a very attractive RAM option with a potentially high payoff. Compared to SFQ RAM, MRAM is non-volatile, non-destructive readout, and higher density, albeit slower. Compared to hybrid JJ-CMOS, monolithic RSFQ-MRAM does not require signal amplification to convert mV-level SFQ input signals to CMOS operating voltage, which should improve relative access and cycle times. If 50psec SMT switching is achieved, memory cycle times should be between 100 and 200ps. A 200ps cycle time MRAM—where RSFQ-based floating point units and vector registers are communicating with multiple 64-bit wide memory chips—translates to a local memory bandwidth of ~40GBytes/s. Thus, cryogenic superconducting-MRAM could well meet the minimum local memory bandwidth requirement of ~2 Bytes/FLOP for a 200-GFLOP processor.

High-density, stand-alone superconducting-MRAM main memory—even with 500 ps to 1 ns cycle times—could be a big win since this would allow the entire computer to be contained within a very small volume inside the cryostat, reducing the communication latency to a bare minimum in comparison with external main memory. Such an arrangement would expose the entire MRAM memory bandwidth to the RSFQ logic, providing RSFQ processor arrays access to both extremely fast local as well as global main memory. Superconductive-MRAM mass memory should also dissipate less power than RT CMOS, since all high-speed communication would be confined to the cryostat.

RSFQ MRAM density should be able to track that of commercial MRAM with some time lag. Monolithic 4 Mb RSFQ-MRAM chips operating at 20 GHz would provide the low-latency RAM needed for 100 GHz RSFQ processors. It will, however, require investment beyond 2010 to bring this high-payoff component to completion.

Understanding and optimizing the cryogenic performance of MRAM technology is important. Fortunately, device physics works in our favor.

- MRAM has a higher MR at low temperatures due to reduced thermal depolarization of the tunneling electrons. A 50% increase in MR from room temperature to 4 K is typical. Higher MR combined with the lower noise inherent at cryogenic operation provides faster read than at room temperature and lower read/write currents.
- Since thermally induced magnetic fluctuations will also decrease at low temperatures, the volume of magnetization to be switched can be reduced, with a concomitant reduction in switching current and power dissipation.
- Resistive metallization, such as Cu used in commercial MRAM chips, can be replaced with loss-less superconducting wiring. This will dramatically reduce drive voltages and total power dissipation, since ohmic losses are at least an order-of-magnitude higher than the dynamic power required to switch the bits. Calculations indicate that the field energy per pulse approximates 60 fJ per mm of conductor, or 240 fJ to write a 64-bit word. Without ohmic losses, the energy required to write a FS-MRAM at 20 GHz is simply 240 fJ times 20 GHz, or 5 mW at room temperature. Reading would require approximately one-third the power or ~1.7 mW. For SMT with 50 ps switching, power per bit reduces to ~10 fJ. If switching currents can be further reduced below 1 mA, as expected at low temperatures, the dynamic power dissipation in the memory would go down even further. A more accurate assessment of total power dissipated in the memory depends on memory arrangement and parasitic loss mechanisms.

MRAM Major Issues and Concerns

Materials Data

A major concern is the paucity of low-temperature experimental data. Data needs to be acquired early on the fundamental properties of candidate magnetic materials and their performance in a memory element. Extrapolations from 300 K to 4 K appear promising for FS-TMR MRAM but are not assured. SMT memory may be more compatible in device resistance and switching speed with RSFQ logic than FS devices, but it is even more embryonic. It should be emphasized, however, that the opportunity afforded by combining MRAM with RSFQ circuits will remain unexplored unless this ultra-high performance computing niche is funded by government investment.

Collaborators

Another issue is establishing collaboration with a suitable organization for device fabrication, and, ultimately, memory product chips. For superconducting-MRAM development, all Cu wiring needs to be replaced by Nb. Compatibility needs to be addressed and process flow determined. For monolithic RSFQ-MRAM, both MRAM and RSFQ fabrication processes will need to be combined, and process compatibility will be more complex. The major organizations and their activities are noted below:

- DARPA accelerated MRAM device development at IBM, Motorola, and NVE Corporation through seed investment in the 90's.
- Since then, Motorola's Semiconductor Sector (now Freescale Semiconductor) has invested several hundred million dollars to build and operate an MRAM fabrication line currently producing 4Mbit chips.
- Freescale is also collaborating with Honeywell to build MRAM radiation-hardened devices.
- IBM may be another potential partner, although the company has maintained a lower level R&D effort to date (\$8-10 M/year) without any commitment to production.
- NVE Corporation continues some level of activity in spintronics.

Flux Trapping

Flux trapping, which has been identified as a generic issue for RSFQ circuits, is of special concern here since MRAM incorporates ferromagnetic films in the chip. The effects should be evaluated and methods developed to avoid this problem.

3.2.4 MEMORY – SUMMATION

The roadmap and Investment plan for cryogenic RAM envisions initiating development of the three RAM options discussed above. Each technology would be monitored and evaluated for progress and possible show stoppers. Approximately halfway through the five-year cycle, the efforts and investment allocation would be critically reviewed.

The expected relative merits of each technology as the panel understands them are given in Table 3.2-1. These would be updated periodically during the time of the roadmap.

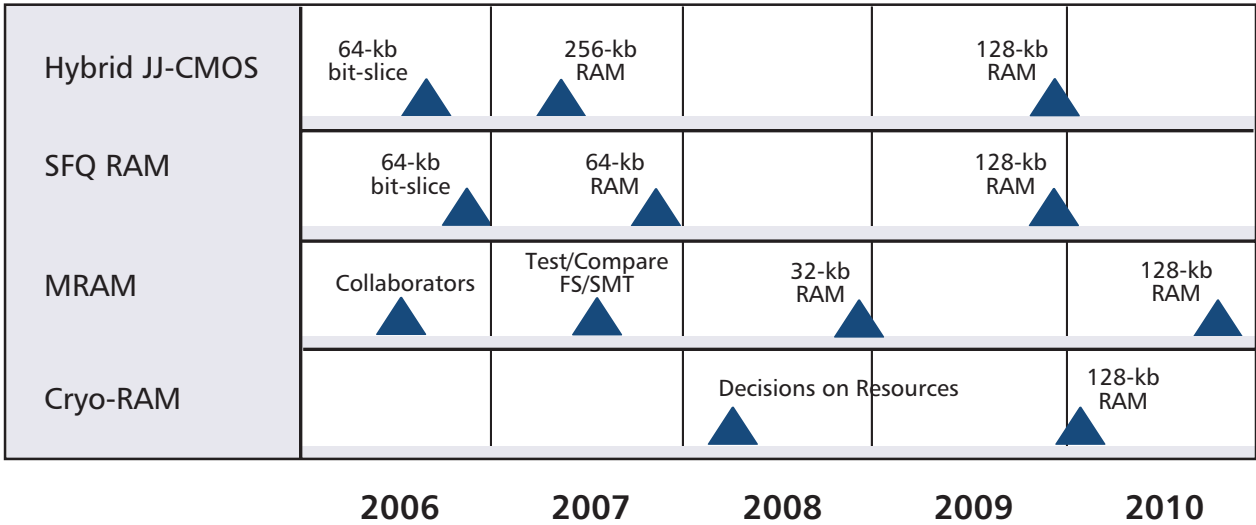


Figure 3.2-4. Cryogenic RAM roadmap.

Memory Investment Summary

TABLE 3.2-10. INVESTMENT PLAN FOR CRYOGENIC RAM (in \$M)				
Memory Technology	JJ-CMOS RAM	SFQ RAM	JJ-MRAM	Memory Total
Investment	\$14.2	\$12.5	\$22	\$48.45

3.3 CAD TOOLS AND DESIGN METHODOLOGIES

A CAD suite that enables a competent engineer who is not a JJ expert to design working chips is essential to demonstrate the readiness of RSFQ processor technology described in Section 3.1. Presently available software is satisfactory only for modest-size circuits of a few thousand JJs, and CAD functions unique to superconductive circuits are poorly supported.

Today the superconductor industry primarily uses the Cadence environment, augmented by custom tools. Stony Brook University has developed custom programs, including a device-level circuit simulator and an extractor of inductance from the physical layout. The University of Rochester has moved superconductive IC design to the Cadence software environment. However, the custom tools are poorly supported by the commercial vendor, if at all, and significant investment will be necessary to integrate the superconductive electronics (SCE) tools to arrive at a maintained CAD suite that will support VLSI-scale circuit design.

CAD capability needs to advance to support an increase in integration scale from a few thousand to millions of JJs:

- Inductance extraction must be extended to sub-micron wires.
- Device and noise models must be extended to sub-micron JJs.
- Transmission line models must be extended to the high frequency regime.
- VHDL models and methods must be extended to larger integration scale and more complex architectures.

Status

In the last decade, superconductive RSFQ digital electronics has grown in size and complexity as it has moved from academic research toward applications development. Initially, only a few experts could successfully design circuits of significant functionality. Two examples of what has been achieved are the line of HYPRES A/D converters, including the 6,000 JJ design pictured in Figure 3.3-1 and the Japanese-developed microprocessor shown in Figure 3.3-2. The technology is now accessible to a wider group of designers, due primarily to the introduction of standard design methodology and industrial CAD tools.




Figure 3.3-1. ADC chip containing about 6,000 JJs was developed by HYPRES. This integration scale can be achieved by experts in the field using design rule checking (DRC) CAD verification.




Figure 3.3-2. A bit-serial superconductor microprocessor featuring 6,300 JJs, 7 instructions, and a 16 GHz clock. The circuit was developed and tested by an independent consortium of Japanese universities.

Readiness

The IC design software developed by Stony Brook University and the University of Rochester was integrated into a single suite at Northrop Grumman. The readiness of U.S.-based design methodology and CAD tools is described below using the Northrop Grumman capability as an example. While built upon the academic projects, it also leverages the methodology and software that serves the commercial semiconductor ASIC world.

A significant challenge for fabrication at or below the 0.25-micron node is the development of inductance extraction software that is accurate for smaller features. Cadence supports extraction of device parameters from the physical layout. However, extraction of inductance values is not well-supported. New software should use 3-dimensional modeling of magnetic fields to accurately predict the inductance of sub micron lines.

Hardware Description Language (HDL) models need to contain the right information, such as gate delay as a function of the parasitic inductive load associated with the physical wiring between gates. Standards for gate models exist, such as standard delay format in which gate delay is given in terms of nominal, typical high, and typical low values of physical parameters.

VHDL has been used in the design of such complex multichip superconductive digital systems such as Northrop Grumman's programmable bandpass A/D converter that combined mixed signal, digital signal processing, and memory elements. The same schematic used to design the circuit in VHDL can be used for layout-versus-schematic (LVS) verification.

By combining process control monitor and visual inspection data to characterize each die on the foundry side with CAD verification on the designers side, first-pass success had become routine for chips of up to a few-thousand JJs fabricated in the Northrop Grumman foundry.

Issues and Concerns

Present simulation, layout, and verification tools could form the foundation of a new CAD capability. Translating existing CAD capability, DRC, LVS, and generation of Pcells for physical layout, to a new foundry process or processes should be quick and relatively easy. Additional CAD development should be driven by advances in feature size, signal frequency, integration scale, and circuit complexity, all of which need to leapfrog present capability by several generations.

Physical models

A mature RSFQ technology will operate closer to ultimate physical limits than present technology. Ultimately, this will require new models of junctions, quantum noise, transmission lines, and inductors for physical simulation. Points for consideration include:

- **JJs.** The present RSJ model may not be adequate for sub-micron, high J_C junctions.
- **Quantum noise.** Quantum noise may add significantly to thermal noise in high J_C junctions.
- **Transmission lines.** These are presently considered dispersion- and attenuation-free, which may not be adequate for signals with frequency content approaching the Nb gap frequency (about 800 GHz).
- **Inductors.** Kinetic inductance becomes increasingly large relative to magnetic inductance at sub-micron linewidths and may be frequency dependent at high frequency.

All of these effects should be manageable if they can be captured in physical simulation. Standard physical simulators have already been extended to include superconductive elements such as JJs (e.g., WRSpice). Addition of new elements may now be required.

Parameter extraction from physical layout is needed both for LVS and for back annotation, whereby a netlist is generated from the physical layout. Verification is presently done without checking inductance values. Back annotation presently can only be done at the gate level using LMeter. A true 3D EM (3-dimensional electromagnetic) algorithm may be required to attain high accuracy at sub-micron sizes. Both the physical-level simulator and the inductance extraction tools should be integrated into the existing CAD environment.

Hardware Description Language

VHDL simulation methods will also require further development. More sophisticated modeling will be required for complex, random logic circuits. Standard delay format could be readily implemented in the usual way, but it is not clear whether this would be effective in superconductive circuits. Also at issue are the effects of signal jitter and of timing-induced probabilistic switching in RSFQ gates. These may combine to make effective hold and setup times in the low BER regime significantly larger than idealized, noiseless circuit simulation would indicate. While the mathematical formalism is well understood, it has not been implemented in CAD for automated checking.

Board design and packaging design could proceed using standard 3D EM software (e.g., HFSS) and standard methods, in which models generated in the frequency domain are translated to the time domain for use in physical-level simulations. Software development may be required because the frequency range of interest extends up to 1 THz, and EM modeling must be able to handle superconductors. This will require modification and extension of the commercial software package.

Roadmap

Table 3.3-1 illustrates the roadmap and milestones for acquisition and development of CAD tools to meet the requirements of the processor and circuits roadmap.

TABLE 3.3-1. ROADMAP AND MILESTONES FOR ACQUISITION AND DEVELOPMENT OF CAD TOOLS							
Task	2005	2006	2007	2008	2009	2010	Milestone
DRC, LVS, PCells							Layout & verification
Place-and-route							Macro circuit designed
Device Models							Simulation & test agree
Induction Extraction							LVS with values
VHDL							Timing design method
Auto Logic Synthesis							Macro circuit designed
Supercon. EM software							Verify in experiment
CAD tool integration							Unified CAD suite
CAD training & support							Designs work first pass

The roadmap indicates that most CAD tasks must be completed early in the program to make them available to circuit designers.

Investment

The investment estimated for the design and CAD development activity is \$24.1 million, as detailed in Table 3.3-2. This includes standard tasks plus the development of device level, logic level, and board level capability.

TABLE 3.3-2. ESTIMATED INVESTMENT FOR DESIGN AND CAD DEVELOPMENT	
Task	Cost (\$M)
Pcells, DRC & LVS Verification	2.0
Place-and-route	1.2
Physical Device Models	4.4
Inductance Extraction	3.4
Cadence Suite Tool Integration	1.3
VHDL Timing Analysis & Design	4.3
Auto Logic Synthesis	2.2
Superconductor 3D EM Software	3.1
CAD Training & Support	2.2
Total Investment	24.1

Present design and CAD capability has been developed with moderate government funding coupled to a comparable measure of industry support. Table 3.3-3 depicts the expected results for three different levels of government investment.

TABLE 3.3-3. EXPECTED IMPROVEMENT IN CAD TOOLS	
Funding Level	Expected Improvement
Full Funding	<ul style="list-style-type: none"> – Build on existing design and CAD tools. – CAD tools will be developed that meet the roadmap schedule and goals.
Moderate Funding	<ul style="list-style-type: none"> – Minimum improvement to meet the needs of HEC. – Slowly and only in those areas where effort is required to achieve an industry result.
No Funding	<ul style="list-style-type: none"> – No improvement in design capability to meet the needs of HEC.

3.4 SUMMARY

The total investment over a five-year period for development of RSFQ processor, cryogenic RAM, and supporting CAD tools is \$118.95 million. Table 3.4-1 shows the allocation of this investment over the three elements.

TABLE 3.4-1. TOTAL INVESTMENT REQUIRED FOR PROCESSORS, MEMORY, AND CAD DEVELOPMENT	
Requirement	Cost (\$M)
RSFQ Processors	46.4
Cryogenic RAM	48.45
CAD Tools and Design Methodologies	24.1
Total Investment	118.95

04

By 2010 production capability for high density IC chips will be achievable by application of manufacturing technologies and methods similar to those used in the semiconductor industry.

Yield and manufacturability of known good die can be established and costs understood.

The 2010 capability can be used to produce chips with speeds of 50 GHz or higher and densities of 1-3 million JJs per cm^2 .

Beyond the 2010 time frame, if development continues, a production capability for chips with speeds of 250 GHz and densities comparable with today's CMOS is achievable.

Total Investment over five-year period: \$125 million.

SUPERCONDUCTIVE CHIP MANUFACTURE

The promise of digital Rapid Single Flux Quantum (RSFQ) circuits is now well established by the results that have come out of various research laboratories. Basic digital device speeds in excess of 750 GHz have been demonstrated. Today, superconductive electronic (SCE) chip manufacturing capability is limited to $10^4 - 10^5$ Josephson junctions (JJs) (equivalent to a few thousand gates) with clocked logic speeds of less than 20 GHz, due to fabrication in an R&D mode with old fabrication equipment and technologies. The use of SCE to achieve petaflops-scale computing will require fabrication of large-area, high-density, superconductive IC chips with reasonable yield. In this section, the panel will:

- Assess the status of IC chip manufacturing for superconductive RSFQ electronics at the end of calendar year 2004.
- Project the capability that could be achieved in the 2010 time-frame.
- Estimate the investment required for the development of RSFQ high-end computers within approximately five years.

The use of SCE to achieve petaflops-scale computing will require fabrication of large-area, high-density, superconductive IC chips with reasonable yield.

Table 4-1 summarizes the roadmap for chip production for aggressive government, moderate, and no funding scenarios.

TABLE 4-1. SCE INTEGRATED CIRCUIT CHIP MANUFACTURE ROADMAP		
Aggressive Funding	<ul style="list-style-type: none"> – Establishment of a modest volume manufacturing capability and infrastructure for delivery of known good die. – Clock rates of 50 GHz or higher, densities of 1-3 million JJs per cm² by 2010. Yield and manufacturability established and costs understood. 	\$125 M
Moderate Funding	<ul style="list-style-type: none"> – Establishment of low volume pilot line/R&D capability with some upgrades to present R&D capabilities. – Pilot/R&D capability demonstrations of density and clock rates of interest by 2014. Yield and manufacturability will not be demonstrated. 	~\$60 M
No Funding	<ul style="list-style-type: none"> – Continued R&D efforts in academic and foreign laboratories. Low continuity of effort. – Modest increases in clock rate and circuit densities in R&D demonstrations. Performance, gate density, and manufacturability inadequate for petaflops computing. 	0

Potential exists for further improvement of the technology beyond 2010. That vision is summarized in Table 4-2. The roadmap established by the panel, detailed in section 4.6, provides for the development of two new generations of RSFQ chip technology by 2010 to be developed in a pilot line and then transferred to a manufacturing line. The first generation process, which is projected for 2007, builds on the most advanced process demonstrated to date, adding minimal improvements but using newer equipment based on 250 nm complementary metal oxide semiconductors (CMOS) processing.

The second generation process, which is projected for 2009, assumes narrower line widths (with the same 250 nm lithographic tools), a modest increase in critical current density, J_c , and the introduction of well understood planarization processes from CMOS.

Beyond 2010, a significantly denser, faster chip technology could be developed, but in a time frame, and requiring resources, outside of the roadmap outlined in this study. This scenario assumes migration to 90 nm or better equipment and processes, moving to layer counts comparable to future CMOS technology, and aggressively increasing the current density to achieve junction speeds at the limit of Nb technology (1,000 GHz).

TABLE 4-2. TECHNOLOGY PROJECTIONS FOR RSFQ PROCESSOR CHIPS

	Units	1st Generation (2007)	2nd Generation (2009)	Vision (beyond 2010)
TECHNOLOGY PROJECTIONS				
Technology Node		0.5-0.8 μm	0.25 μm	90 nm or better
Current Density	kA/cm ²	20	50	>100
Superconducting Layers	Count	5	7-9	>20
New Process Elements		None	Planarization	Alternate barriers, additional trilayers, vertical resistors, vertical inductors, etc.
Power		Modest Improvements	Reduced Bias Voltage	CMOS-like [AWK2], lower I_c
PROJECTED CHIP CHARACTERISTICS				
JJ Density	MJJs/cm ²	1	2-5	250
Gate Clock Rate	GHz	25	50-100	250
Power per Gate	nW/GHz/gate	15	8	0.4

JJ density and clock rates for the roadmap were estimated based on scaling of existing designs with the improvements in line pitch and layout taken into account. The vision of performance beyond 2010 is an estimate based on the assumption that some number of improvements will allow JJ densities comparable to CMOS 90-130 nm node circuitry and innovative circuit designs will be developed to capitalize on these improvements. If multiple JJ layers and extensive vertical structures for components—such as inductors and resistors—that are now in-plane are introduced, these densities are feasible. Since some of the density improvement may require using additional active devices where passive ones are now used, the number of JJs per gate may increase. Clock rate is estimated to be a factor of four lower than the maximum JJ speed (1,000 GHz). This requires significant advances in design technology in order to beat the current factor of 6-9 derating of the clock. Power for the first generation process was calculated with only modest, known improvements to the biasing methods currently used.

The second generation requires some advancement in design technology to find ways to compensate for a modest 10% drop in margins with reduced bias voltage (e.g., without design improvements a 35% gate margin will drop to 25%). The vision beyond 2010 assumes significant improvement in design technology, with implementation of a more CMOS-like biasing structure, such as the Self Clocked Complementary Logic (SCCL) proposed by Silver and Herr, and aggressive scaling of critical current, I_c , to reduce intrinsic power.

4.1 SCE IC CHIP MANUFACTURING – SCOPE

The primary function of the SCE IC chip foundry is to develop and maintain a standard fabrication process with device and circuit density suitable for petaflops chips and to provide manufacturing capability with throughput and yield of functional chips sufficient to construct a petaflops-scale machine. It is estimated that the baseline first generation process meeting these requirements would include 5-7 Nb layers, and 10^6 , 0.8 μm -diameter JJs per IC chip. The JJ critical current density (J_c) would be 20 kA/cm². The panel anticipates a demonstrated manufacturing capability of at least 10^4 functional chips per year will be required to provide the IC chip manufacturing for a petaflops-scale computer following the capability demonstrations outlined in this study.

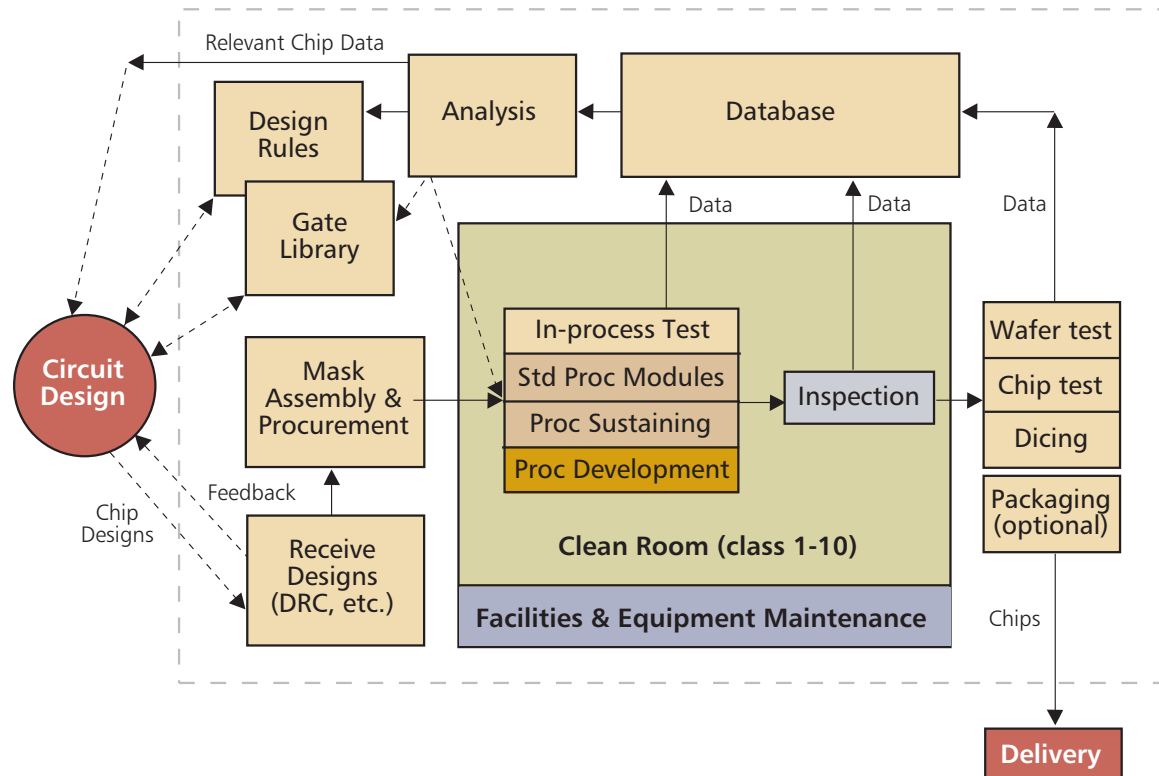


Figure 4-1. Schematic diagram of functions included in IC chip manufacture (shown inside the dashed box).

The functions included in IC chip manufacture are shown schematically in Figure 4-1.

Primary fabrication activities occur within the clean room facility. These include:

- Process development.
- Sustaining the manufacturing process and standard process modules.
- Facilities and equipment maintenance.
- Inspection.
- In-process testing at room temperature.

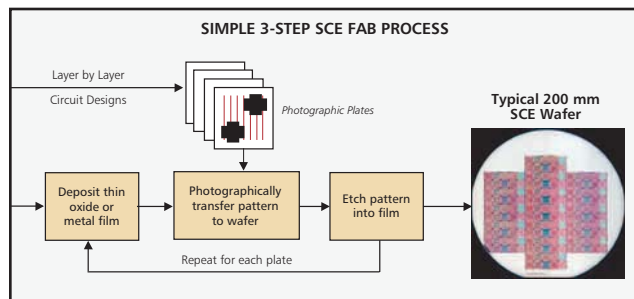
All IC chips on completed wafers undergo room temperature testing prior to dicing. Process control monitor (PCM) chips are mounted on superconductive multi-chip modules (MCMs) tested at 4 K. These Nb-based MCMs are produced in the main fabrication facility, while MCMs for product IC chips are produced separately. (MCM fabrication and testing is discussed in Chapter 6.)

The fabrication group:

- Maintains a database of PCM data and is responsible for its analysis, which provides the basis for the process design rules.
- Maintains the gate library, which is developed in collaboration with the circuit design group(s); is responsible for incorporating circuit designs into mask layouts and for procuring photolithographic masks.
- Produces IC chips for verification of CAD, gates, designs, IC chip architectures, and all IC chips required for the demonstrations associated with this program.
- Potentially can provide foundry services for other digital SCE programs (government, academic, or commercial).

4.2 DIGITAL SUPERCONDUCTIVE IC CHIP FABRICATION – STATUS

Fabrication of superconductive ICs is based on semiconductor wafer processing, using tools developed for CMOS processing or customized versions of these tools. However, the process flow is much simplified from that of CMOS, with only metalization or insulator deposition steps, followed by patterning of these layers by etching. Figure 4-2 shows a simplified flow, along with a photo of a superconducting foundry (one tunnel).



a) Simple process flow for fabrication of RSFQ circuits



b) RSFQ fabrication facility

Figure 4-2. Fabrication of RSFQ ICs is accomplished using semiconductor equipment and processes.

Significant activity in the area of digital superconductive electronics has long existed in the United States, Europe, and Japan. Over the past 15 years, Nb-based integrated circuit fabrication has achieved a high level of complexity and maturity, driven largely by the promise of ultra-high-speed and ultra-low-power digital logic circuits. An advanced process has one (JJ) layer, four superconducting metal layers, three or four dielectric layers, one or more resistor layers, and a minimum feature size of $\sim 1\ \mu\text{m}$. Today's best superconductive integrated circuit processes are capable of producing digital logic IC chips with 10^5 JJs/cm². Recent advances in process technology have come from both industrial foundries and university research efforts, resulting in reduced critical current spreads and increased circuit speed, circuit density, and yield. On-chip clock speeds of 60 GHz for complex digital logic and 750 GHz for a static divider (toggle flip-flop) have been demonstrated. Large digital IC chips, with JJ counts exceeding 60,000, have been fabricated with advanced foundry processes. IC chip yield is limited by defect density rather than by parameter spreads. At present, integration levels are limited by wiring and interconnect density rather than by junction density, making the addition of more wiring layers key to the future development of this technology.

Recent advances in process technology have come from both industrial foundries and university research efforts.

Nb-based superconductive IC chip fabrication has advanced at the rate of about one generation, with a doubling of J_c , every two years since 1998. Increasing J_c enables increasing digital circuit speed. This is illustrated in Figure 4-3, a plot of static divider speed from several of sources versus J_c . Points beyond 8 kA/cm² represent single experimental fabrication runs, not optimized processes. Theoretically, divider speed should approach 1,000 GHz (1 THz), however, the process and layout must be optimized to reduce self-heating effects in junctions with J_c beyond 100 kA/cm². The first generation Nb process discussed in Section 4.6 should be based on 20 kA/cm², 0.8 μm diameter junctions with six superconducting metal layers. Static dividers fabricated in this process have achieved speeds of 450 GHz, which should enable complex RSFQ circuits with on-chip clock rates of 50 to 100 GHz.




Figure 4-3. Demonstrations of RSFQ circuit speed with increasing J_c

Table 4-3 lists the major active institutions engaged in Nb RSFQ circuit production today. Process complexity can be assessed from the number of masking and superconducting wiring layers. J_c , minimum feature size, and minimum JJ area are also listed. Not included are laboratories and companies producing SQUIDs, electromagnetic sensors, and research devices.

Some of the organizations in Table 4-3 provide foundry services to industrial, government, and academic institutions. HYPRES, Inc. has tailored its 1-kA/cm² fabrication process to allow a large variety of different IC chip designs on a single 150-mm wafer. Its 2.5 and 4.5 kA/cm² IC chips represent the current state of the art in the United States. Foundry services from SRL enable many groups in Japan to design and test new circuit concepts. The SRL foundry in Japan offers a 2.5-kA/cm² process and is developing a 10-kA/cm² process. Northrop Grumman Space Technology closed its foundry in mid-2004. Prior to that, its 8 kA/cm² process, which was being upgraded to 20 kA/cm² in 2004, represented the state of the art in the world. An excellent survey paper on the state of the art in fabrication technology appeared in the *Transactions of the IEEE*, October 2004, and is reproduced as Appendix I: Superconductor Integrated Circuit Fabrication Technology. (The full text of this appendix can be found on the CD accompanying this report.)

TABLE 4-3. REPRESENTATIVE NB IC PROCESSES						
Institution	No Masks	J_c (kA/cm ²)	Min. JJ Area (μm^2)	Min. Feature (μm)	Wire Layers	Primary Applications
NGST	14	8*	1.2	1.0	3	RSFQ logic, ADC
HYPRES (4.5kA)	11	4.5, 6.5	2.25	1.0	3	RSFQ logic, ADC
ISTEC (SRL)	9	2.5**	4.0	1.0	2	RSFQ logic
HYPRES (1.0kA)	10	1, 2.5	9.0	2.0	3	RSFQ logic, ADC
AIST	7	1.6	7.8	1.5	2	RSFQ logic
IPHT Jena	12	1	12.5	2.5	2	RSFQ logic
Lincoln Lab	8	0.5, 10	0.5	0.7	2	R&D
Stony Brook Univ.	8	0.2 to 12	0.06	0.25	2	R&D
PTB	8	1	12	2.5	1	RSFQ logic
UC Berkeley	10	10	1.0	1.0	2	RSFQ logic
Univ. Karlsruhe	7 or 8	1 to 4	4.0	1.0	2	Analog, RSFQ logic

*NGST was introducing a 20 kA/cm² process when its foundry was closed in 2004.

**SRL is upgrading to a 10 kA/cm² process with 6 Nb layers.

4.3 SCE CHIP FABRICATION FOR HEC – READINESS

Implementation of SCE for petaflops computing will require three circuit types: logic, cryogenic memory, and network switches. For the HTMT petaflops project, it was estimated that 4,096 processors, comprised of 37K SCE IC chips containing a total of 100 billion JJs and partitioned on 512 MCMs, would be required. Manufacture of such a large-scale system will require significant advances in SCE IC chip fabrication.

The two key advantages of superconductive RSFQ technology for petaflops-scale computing are ultra-high on-chip clock speeds (50-100 GHz or greater) and ultra-low power dissipation (nanowatts per gate). However, while the status presented in the previous section shows the enormous promise of RSFQ technology, significant improvements in density and clock speed of chips are required for petaflops computing. In order to produce petaflops processor elements with a practical chip count per element (~10-15), the RSFQ logic chip will have an estimated 12 million, 0.5 - 0.8 μm diameter junctions and 5 - 7 wiring layers on a die no larger than 2 cm x 2 cm. Junction current density must be increased to provide higher circuit speeds (circuit speed increases as the square root of the J_c increase). Scaling present-day Nb IC chip technology to produce the necessary circuits will require:

- A major increase in circuit density (from <100K to 3M JJs/cm²).
- A decrease in minimum feature size (from ~1 μm to ~0.25 μm).
- An increase in J_c (from 2 – 4.5 kA/cm² to at least 20 kA/cm²).
- The addition of several superconducting interconnect layers.
- 2,000 – 5,000 pin-outs per IC chip.

Two key advantages of superconductive RSFQ technology for petaflops-scale computing are ultrahigh on-chip clock speeds and ultralow power dissipation.

While this list represents significant advances in current JJ fabrication technology, the chip processing requirements are not challenging when compared with CMOS processing. The advanced RSFQ IC chip process requires the development of:

- A robust 0.5 - 0.8 μm junction technology.
- Small-footprint resistors.
- Planarization.
- Associated plug technology for vias.

The process development listed represents application of tools, techniques, and methods that were current in CMOS in the mid to late 90's. This means that the development needed by RSFQ can build on well understood technology. Table 4-4 summarizes the readiness of RSFQ chip fabrication technology for development to the levels required for petaflops by 2010. The issues and concerns listed are discussed in detail in Section 4.5.

TABLE 4-4. READINESS OF RSFQ CHIP FABRICATION TECHNOLOGY FOR PETAFLOPS			
Element	Readiness	Requirement	Issues/Concerns
Circuit Speed	High	2X-5X Increase in J_C	Degradation of yield due to higher parametric spreads
Chip Density	High	1995-98 CMOS process technology	Higher device parametric spreads
Yield	Moderate	Greater than 20-30% for required number of chips	Existing yields not quantified
Known-good-die production	Moderate	Cryogenic testing for delivery	Cost and time for full testing

4.4 IC CHIP MANUFACTURING CAPABILITY – PROJECTIONS

The recent evolution of superconductive IC chip fabrication technology and its extension to 2009 and beyond is outlined in Table 4-5. The first three columns represent the NGST foundry process through early 2004. The next column represents a 0.8 μm , 20 kA/cm² process that was demonstrated by, and being adopted at, NGST when its foundry closed in 2004. The final two columns represent the processes anticipated for the proposed program, with petaflops-scale computing achievable with the technologies introduced in the years 2007 and 2009.

TABLE 4-5. SUPERCONDUCTIVE IC CHIP TECHNOLOGY ROADMAP						
Year	1998	2000	2002	2004	2007	2009
Minimum Feature size (μm)	1.5	1.0	1.0	0.60	0.60	0.25
Junction size (μm)	2.50	1.75	1.25	0.80	0.80	0.50
Junction current density (kA/cm ²)	2	4	8	20	20	50
Wafer diameter (mm)	100	100	100	100	200	200
Superconducting Wiring Layers	4	4	4	4	5	7
Planarization	no	no	no	partial	partial	yes
Clean Room class	100	100	10-100	10-100	10-100	1-10
Wafer starts per year	12	200	200	300	1K	10K

The superconductive IC chip fabrication process will build upon existing experience and concepts that already have been proposed or demonstrated in superconductive circuits. Figure 4-4 is a cross-sectional view of a conceptual superconductive IC chip process corresponding to 2009 in Table 4-5. It relies extensively on metal and oxide chemical-mechanical planarization (CMP). This process has one ground plane, four wiring layers (including base electrode), two resistor layers, self-aligned junction contacts, and vertical plugs or pillars for interconnection vias between wiring layers. The aspect ratio of the vertical plugs is on the order of 1:1, which does not require the complex chemical vapor deposition or hot metal deposition processes typically used in semiconductor fabrication to fill vias of more extreme aspect ratio. Power lines and biasing resistors are located below the ground plane to isolate the junctions from the effect of magnetic fields and to increase circuit density. The final superconductive IC chip process is likely to include an additional wiring layer and/or ground plane and vertical resistors to reduce the chip real estate occupied by junction shunt resistors.

Superconductive IC chip fabrication for petaflops-scale computing far exceeds the capability of existing superconductive foundries, even the former NGST foundry. However, the required level of integration has been available for several years in the semiconductor industry. Comparison with the Semiconductor Industry Association (SIA) roadmap suggests that a Nb superconductive process implemented in 1995 CMOS technology would be adequate. The fabrication tools, including advanced lithography, CMP, and infrastructure support, are readily available, so no major new technologies or tools are required to produce the superconductive IC chips needed for petaflops-scale computing.

4.5 IC CHIP MANUFACTURE – ISSUES AND CONCERNS

RSFQ logic offers an extremely attractive high-speed and low-power computing solution. Low power is important, because it enables both high IC chip packaging density and low-latency interconnects. For large systems, total power is low despite the penalty associated with cooling to cryogenic temperatures.

The major issues relevant to the circuit complexity, speed, and low power necessary for petaflops-scale computing are discussed below.

4.5.1 IC CHIP MANUFACTURE – IC MANUFACTURABILITY

In section 4.3, the panel listed the technical challenges involved in the fabrication of ICs suitable for petaflops-scale computing. SCE technology will have to improve in at least three key areas including reduction in feature size, increase in layer count, and increase in gate density. While the required level of complexity is commonplace in high-volume semiconductor fabrication plants, it represents a significant challenge, particularly in terms of yield and throughput.

Scaling to petaflops requires more than a 10-fold increase in circuit density along with a decrease in feature sizes to several times smaller than present day technology.

At present, Nb-based JJ technology is at the $10^4 - 10^5$ JJ/cm² integration level, with minimum feature sizes of ~ 1 μ m, depending on the foundry. Although the technology can support up to 10^5 JJ/cm², circuits of this scale have not been demonstrated due to other yield-limiting factors. Scaling to petaflops requires more than a 10-fold increase in circuit density to at least 10^6 JJ/cm², along with a decrease in feature sizes to several times smaller than present day technology (to 0.5 - 0.8 μ m for junctions and ~ 0.25 μ m for minimum features), as shown in the technology roadmap (Table 4-5). Additional superconducting interconnect layers are also required.

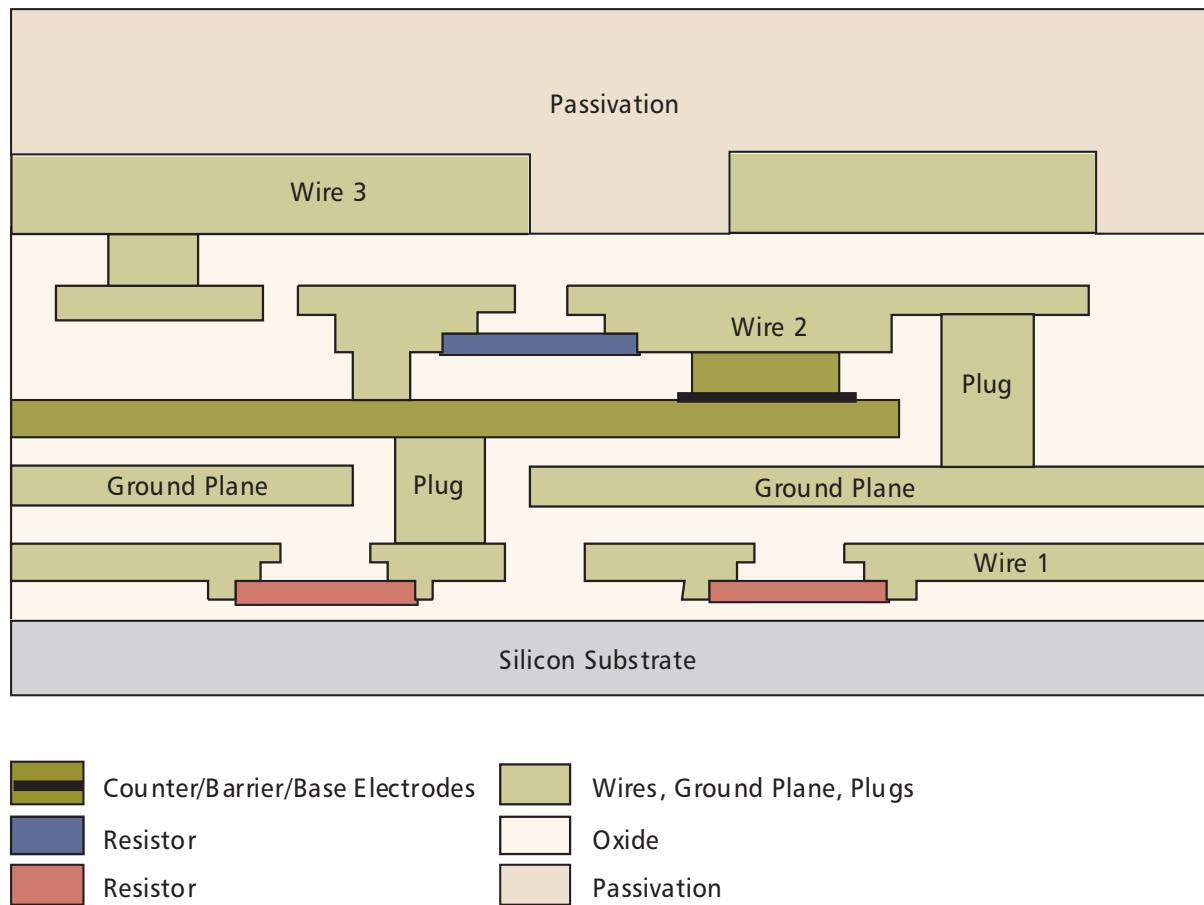


Figure 4-4. Notional cross-section of the superconductive IC chip fabrication process illustrating salient features of an advanced process, which includes four interconnect levels, planarization, and plugs.

The leap to gate densities and feature sizes needed for petaflops cannot be achieved in a single step. Sequential improvements will be needed, commensurate with the funding available for the process tools and facilities. Semiconductor technology has scaled by a factor of 0.7x in feature size and 2.5x in gate density for each generation, each of which has taken about three years and a substantial capital investment. However, much of the delay in advancing semiconductor IC chip technology was due to the lack of advanced process tools, especially in the area of photolithography. The panel expects to achieve 0.8 μm junction sizes in the next generation, because the tools and methodologies are already available. In fact, NGST, in collaboration with JPL, demonstrated a 0.8 μm junction-based circuit fabrication process and was on the way to adopting it as standard when its Nb foundry closed in 2004. Further reduction in feature size and other process improvements will occur in subsequent generations.

4.5.2 IC CHIP MANUFACTURE – DEVICE AND CIRCUIT SPEED

RSFQ gate speed is ultimately limited by the temporal width of the quantized pulses. For junctions with critical current I_c and resistance R , this width is proportional to $\Phi_0/I_c R$, where $\Phi_0=2.07\text{mV}\cdot\text{ps}$ is the superconducting flux quantum. The maximum operating frequency of the simplest RSFQ gate, a static divider, is $f_0 = I_c R/\Phi_0$. Maximum speed requires near-critical junction damping, with $2\pi f_0 R C \approx 1$, where C is the junction capacitance. Then $f_0 = (J_c/2\pi\Phi_0 C')^{1/2}$, where J_c and C' are the critical current density and specific capacitance, respectively. These parameters depend only on the thickness of the tunnel barrier. Because C' varies only weakly with barrier thickness, while J_c varies exponentially, f_0 is approximately proportional to $J_c^{1/2}$.

Table 4-6 shows how gate speed depends on critical current density and JJ size. The last three columns list ranges of operating frequency.

TABLE 4-6. DEVICE PARAMETERS AND GATE SPEED				
J_c (kA/cm ²)	JJ size (μ m)	f_0 (GHz)	asynch ckt f (GHz)	clocked ckt f (GHz)
1	3.5	125	52-78	13-31
2	2.5	170	71-107	18-43
4	1.75	230	97-146	24-58
8	1.25	315	132-198	33-79
20	0.8	470	197-296	49-116
50	0.5	700	293-440	73-176
100	0.35	940	394-592	99-232

For a typical Nb RSFQ fabrication process available today, $J_c = 2$ kA/cm², and $f_0 \approx 170$ GHz. However, for both superconductive and semiconductor digital technologies, the maximum clock speed of VLSI circuits is well below f_0 , the speed of an isolated gate. The input and output signals for a logic gate are generally aperiodic, so the SFQ pulses must be well-separated. Since RSFQ gates are generally clocked, timing margins must be built in to ensure minimum clock-to-data and data-to-clock separation. Several timing methods in large, complex circuits—including concurrent or counterflow clocking, dual rail, and event-driven logic—are under investigation. The last two columns list clock speed ranges for asynchronous circuits (such as shift registers) and complex, clocked circuits (ranging from adders and multipliers to processor IC chips, respectively). For example, experience, both experimental and theoretical, shows that complex circuits in 2 kA/cm² technology can be clocked at 20 – 40 GHz. Scaling J_c up to 20 kA/cm² results in $f_0 \approx 450$ GHz, and the panel anticipates that complex circuits built with this technology will be able to operate at speeds ranging from 50 to over 100 GHz. Figure 4-5 is a graphical illustration of these projections.

4.5.3 IC CHIP MANUFACTURE – CIRCUIT DENSITY AND CLOCK SPEED

Successive superconductive IC chip fabrication process generations will result in improved RSFQ circuit performance, with clock rates of at least 100 GHz possible. Because signals propagate on superconductive microstrip lines at one-third the speed of light in vacuum, interconnect latency is comparable to gate delay in the 100 GHz regime. Signal propagation distances must be minimized by increasing gate density, using both narrower line pitch and additional superconductive wiring levels. This may be even more important to increasing the maximum clock rate than smaller junctions and higher critical current density. Thus, successive IC generations focus on improved gate *density* (by reducing line pitch and increased number of interconnect layers) as well as gate *delay* (by reducing junction size). These process improvements are listed and evaluated in Table 4-7.




Figure 4-5. Projections of RSFQ circuit speeds with increasing J_c

TABLE 4-7. ADVANCED PROCESSES ALLOW LATENCY REDUCTION		
	Increased Clock Rate	Increased Density
Process Improvement	<ul style="list-style-type: none"> – Smaller junctions with higher critical current density. 	<ul style="list-style-type: none"> – Smaller line pitch. – Increased vertical integration.
Benefits	<ul style="list-style-type: none"> – Higher circuit speed. – Higher junction impedance. – Higher voltage signals. 	<ul style="list-style-type: none"> – More gates/cm². – Reduced latency* on-chip.
Disadvantages	<ul style="list-style-type: none"> – Larger electrical spreads. – Increased latency* within system. 	<ul style="list-style-type: none"> – Lower yield.

* Latency is measured as the number of clock ticks for signal propagation between a given pair of gates.

4.5.4 IC CHIP MANUFACTURE – PARAMETER SPREADS AND IC CHIP YIELD

Yield is an important factor in manufacturing cost. IC chip yield depends upon many factors. These include:

- Wafer process yield (which depends on throughput).
- Number of masking levels.
- Defect density.
- IC chip size.
- Yield due to parameter variations.

Parameter spreads in superconductive circuits can cause either hard failures or soft failures. Extensive data for important circuit elements such as junctions, resistors, and inductors in a present day superconductive process indicates ~1 - 4% spreads (1σ) for local variations (on-chip), 5% for global variation (across-wafer) and 10% for run-to-run reproducibility. The most critical circuit element is the JJ. Table 4-8 shows on-chip I_c spreads for several Nb IC generations at NGST. (Note that the 2004 column reports early results for a new process that was far from optimized.)

TABLE 4-8. JUNCTION CRITICAL CURRENT VARIATION

Year	1998	2000	2002	2004
Junction size (μm)	2.5	1.75	1.25	0.80
Junction current density (kA/cm^2)	2	4	8	20
ΔI_c (1σ) %	± 1.1	± 1.4	± 1.4	± 2.30
ΔI_c (max-min) %	± 2.5	± 3.4	± 3.5	± 5.9

Junction scaling is necessary to achieve higher clock speeds. It is estimated that 0.5 - 0.8 μm junctions will be required to meet the 50 - 100 GHz clock requirement. Unlike transistors, for which the gate length is the critical dimension (CD) and the circuit is not as sensitive to local variations in gate length, tight areal CD control is required for JJs, since critical current (which scales with junction area) is the important parameter. As feature sizes decrease, absolute CD control will need to improve proportionately. Although the preliminary 2004 result in Table 4-8, obtained at NGST in collaboration with JPL, is promising in that regard, there are few data available on submicron JJ circuits. However, additional data are available for magnetic random access memory (MRAM) technology, which is also based on tunnel junctions with Al oxide barriers, the same barrier used in Nb JJs. IBM has demonstrated a 16 Mb MRAM chip based on deep-sub- μm (0.14 μm^2) tunnel junctions which exhibit resistance spreads of ~2% (1σ). This indicates that 2% J_c control in JJs of similar size is possible.

Local spreads are the most important in determining circuit size, while global and run-to-run variations limit yield (i.e., the number of good chips). Present day spreads are consistent with circuit densities of 1M JJ/ cm^2 , even while taking the low bit-error-rate (BER) requirement of general computing into account (but excluding other yield-limiting effects such as defect density). However, present day tools and methods may not be adequate, and the ability to control CD variation could limit progress towards integration levels beyond 1M JJ/ cm^2 . Commercially available lithography tools provide resolution control of 0.03 μm (2σ) for a feature size of 0.65 μm . For 0.8 μm junctions, this translates to $\pm 5\%$ I_c spread (1σ) for two neighboring junctions, for just the exposure portion of the process. The final CD is a result of several other processes including developing and etching, each of which has a CD tolerance. A net CD tolerance of $\pm 5\%$ may limit yield of large circuits. This suggests that other methods or more advanced tools may be needed.

*Understanding and predicting yield will be important
because it will have a strong bearing on sizing the resources
needed to produce a petaflops-scale computer.*

The extensive modeling of defect density and its effect on die yield developed for the semiconductor industry is directly applicable to superconductive IC chip manufacture because the fabrication processes are similar. Understanding and predicting yield will be important, because it will have a strong bearing on sizing the resources needed to produce a petaflops-scale computer. Defect density determines practical IC chip size. A typical estimate for present defect density in a clean superconductor foundry is 1-2 defects/cm². Defect density on the order of 0.1 defect/cm² will be needed to produce 2 x 2 cm² chips with yields greater than 50%. The optimal IC chip size is a tradeoff between IC chip yield and packaging complexity. The use of Class 1 and 10 facilities and low-particle tools, and good discipline by personnel is required to reduce the process-induced and environment-induced defects to this level.

4.5.5 IC CHIP MANUFACTURE – PRODUCTION OF KNOWN GOOD DIE

Integrated circuit manufacture may be defined as the production of known good die. Producing known good die will depend on the overall circuit yield, which is the product of the individual yields of all the processes used to produce and test the IC chips. For a large-scale system, the number of IC chips required is large (~40,000) and the yield may be low, at least initially. A large volume of wafers will have to be fabricated and tested. Wafer-level, high-speed cryogenic (4 Kelvin) probe stations will be required in order to screen die before dicing. Built-in self-test (BIST) will certainly simplify the screening process. Tradeoffs between the quality of and overhead imposed by BIST should be considered early in the design cycle. Sophisticated BIST techniques have been developed in the semiconductor industry and are readily applied to SCE chips. Simple BIST techniques have already been used to demonstrate high-speed circuit operation on-chip, using a low-speed external interface.

High-speed probe stations and BIST techniques will not be available in the early stages of development. Screening for chips likely to be functional will be limited to low temperature screening of wafers via Process Control Monitor chips and room temperature electrical probing and visual inspection.

For a large-scale system, the number of IC chips required is large (~40,000), and the yield may be low, at least initially.

Experience has shown that IC chip yield in high-volume silicon foundries is low during the initial or pre-production phase, increases rapidly as production increases, and then levels off at product maturity. Due to the volume of IC chips for the petaflops computing, the superconductor foundry will probably operate on the boundary between pre-production and early production phases. Production-type tools will be needed to produce the required volume of IC chips, and larger wafers will be needed to produce the IC chips in a cost-effective and timely way when the number of IC chips is large and the yield is relatively low. The panel projects that 200 mm wafers will be adequate.

Table 4-9 shows case examples of throughput requirements for various yield values, assuming a 24-month production run after an initial development phase to reach 5% yield. As shown, the worst case scenario shows the facility must be sized to produce 1,000 wafers per month. As the product matures, the panel expects to reach well above 20% yield, which will comfortably satisfy the production volume needs. Minimizing the number of different IC chip types (e.g., memory vs. processor) will also reduce the total number of wafers needed.

TABLE 4-9. SCALING REQUIREMENTS FOR SCE IC CHIP PRODUCTION

Yield =	5%	10%	20%	30%
IC chip total	737,280	368,640	184,320	122,880
Wafer total	23,040	11,520	5,760	3,840
Wafer starts per month	960	480	240	160
Chip tests/month @ 4K	30,720	15,360	7,680	5,120

4.6 ROADMAP AND FACILITIES STRATEGY

4.6.1 ROADMAP AND FACILITIES STRATEGY – ROADMAP

The roadmap to an SCE IC chip manufacturing capability must be constructed to meet the following criteria:

- Earliest possible availability of IC chips for micro-architecture, CAD, and circuit design development efforts. These IC chips must be fabricated in a process sufficiently advanced to have reliable legacy to the final manufacturing process.
- Firm demonstration of yield and manufacturing technology that can support the volume and cost targets for delivery of functional chips for all superconductive IC chip types comprising a petascale system.
- Support for delivery of ancillary superconductive thin film technologies such as flip-chip bonding, MCM's, and board-level packaging for technology demonstrations.
- Availability of foundry services to the superconductive R&D community, and ultimately for other commercial applications in telecommunications, instrumentations, and other applications.

From these broad criteria, a few “rules of the roadmap” can be derived. First, rapid establishment of an advanced process with minimal development is desirable. Examples of pilot and R&D level processes of this type are the NGST 20 kA/cm² process and the NEC 10 kA/cm² process. However, such an early process is not expected to be sufficient to meet all the needs of a petaflops-scale system, so development and process upgrades must be planned as well. Early establishment of such a pilot process will not allow rigorous establishment of manufacturing facilities, equipment, and processes, so this must also be planned in separately. In a more relaxed scenario, establishment of a manufacturing capability could be in series with the pilot capability, but the relatively short five-year time frame contemplated does not allow this. In order to assure the most cost effective in introduction of manufacturing capability, use of the Intel “Copy Exactly” process development method should be adopted to the greatest extent possible.

A second possibility is development and establishment of capability at multiple sites or sources. Unfortunately, the development of each generation of process at different sites—perhaps from different sources—does not allow implementation of any form of the Intel “Copy Exactly” philosophy and greatly increases the cost and schedule risk for the final manufacturing process.

The panel believes that the most cost effective five-year roadmap is development and establishment of the SCE chip processes at a single site, perhaps even a single facility originally sized for the final process capability but initially facilitized for the early pilot process requirements. Consideration for copying the manufacturing process exactly to a second source once the volume and/or business requires it, should be designed into the contract or other legal charter establishing the first facility. Finally, while establishment of the IC chip manufacturing capability can be accomplished at a government, academic, or industrial location, it is desirable that this facility to provide for both the future commercial business of supplying petaflops-scale computing system manufacturers with IC chips and supplying other commercial and R&D applications. This would require a strong industrial component, with government oversight derived from the contractual arrangement established for development. The most likely structure is to have a partially, or wholly, government-owned facility, managed and manned by a contractor, with commercial uses offsetting the government's costs. Figure 4.6 shows a roadmap that takes all of these elements into account.

The most cost-effective 5-year roadmap is to develop and establish the SCE chip process at a single facility partially or wholly government owned but managed and manned by a contractor.

As Figure 4-6 illustrates, there are three major elements to the establishment of a superconductive IC chip manufacturing capability for petaflops-scale computing. The first is the establishment of an early pilot line capability in the most aggressive process that still allows early availability of IC chips. Until the first wafer lots are available in this process, development of concepts in CAD, micro-architecture, and circuit design could be accomplished by gaining access to the NEC process, which is the most advanced in the world .

Second, in parallel with the pilot line, facilities and equipment for higher volume production in more aggressive design rules need to be started, as this element will take a substantial lead time before it is available to a program. Once both the pilot line and the manufacturing line are established, processes developed in the pilot line will be transitioned to the manufacturing line. As newer processes become well established, older processes will be phased out, first in the pilot line, then in the manufacturing line. A set of parametric and yield evaluation vehicle IC chips will be developed and held in common for the two lines.

In the model shown in Figure 4-6, all development occurs in the pilot line. If co-located, the two lines could share certain equipment (such as lithography) but would require separate equipment for elements where changing process conditions might jeopardize yield. Using the cluster tool mode of fabrication would minimize redundancy and costs and allow flexibility in assignment of pilot versus manufacturing roles.

The third element is the required support for the IC chip processing. Its two major elements are:

- Testing and screening of parametric and other process control IC chips, as well as yield and functional testing.
- Packaging of IC chips for delivery to the design development and demonstrations of capability.

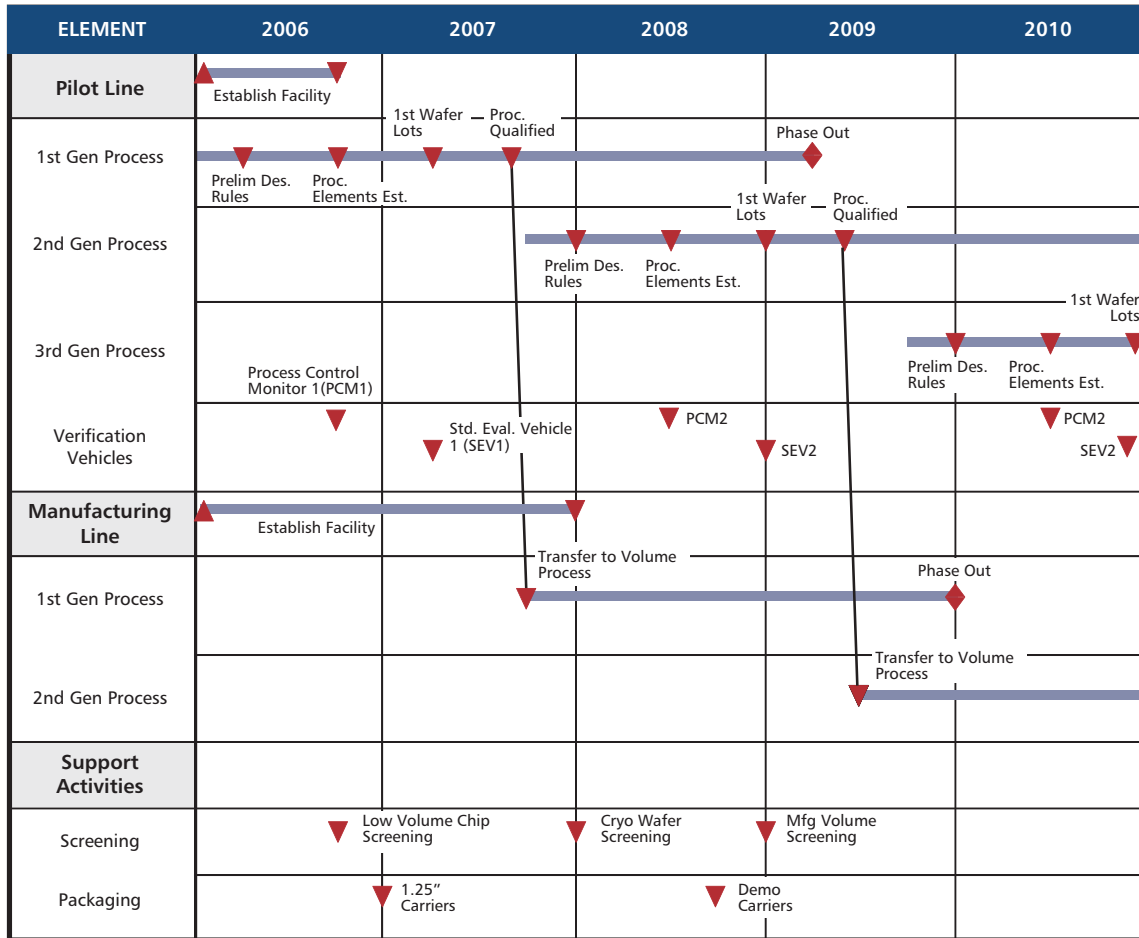


Figure 4-6. Timeline for development of SCE manufacturing capability.

4.6.2 ROADMAP AND FACILITIES STRATEGY – MANUFACTURING FACILITIES

Manufacturing RSFQ IC chips of the required complexity and in the required volumes for petaflops-scale computing will require investment, both recurring and nonrecurring. The recurring costs are associated with operation of the fabrication facility, tool maintenance, and accommodation of new tools. These include the costs of staffing the facility. The nonrecurring costs are mainly associated with the procurement cost of the fabrication tools and one-time facilities upgrades.

Comparison with the SIA roadmap suggests that 1992 semiconductor technology is comparable to that needed for petaflops. Advanced lithography tools, such as deep-uv steppers, will provide easy transition to sub-micron feature sizes with excellent alignment, resolution, and size control. CMP for both oxides and metals will be important as the number of interconnects increase and feature sizes decrease. Timely procurement of the tools will be necessary to achieve the technology goals in the time frame indicated.

At present, there are no true production lines for SCE chips anywhere in the world. In discussing the required manufacturing facilities, the panel assumed a baseline process with 0.8 μm diameter, 20 kA/cm², 5 Nb wiring layers, and 1,000,000 JJs/chip. In such a facility, SCE IC chip production is expected to fall between 10,000 and 30,000 KGD per year.

Process stability is key to the success of a project of this scale. In contrast to typical practice in R&D superconductive IC chip fabrication, the fabrication process on the manufacturing line should not be modified once it is established. In general, practices of the semiconductor industry should be adopted, including:

- State of the art facilities (e.g., class 1-10 clean room)
 - Size and build adequate clean room floor space.
- Standard production tools
 - Standardize on 200 mm wafers.
 - Additional process tools as needed to meet demand Deep-UV stepper Facilities to support CMP production tools (both oxide and metal CMP maybe required).
 - Automated process tools wherever possible.
 - Minimal handling of wafers by human operators.
- Statistical process control and electronic lot travelers and tracking.
- Assemble an experienced fab team, with perhaps 75% of personnel from the semiconductor industry and 25% with extensive superconductor experience.

Semiconductor industry experience indicates that initial IC chip yields will be low (<10%), but will increase as the process matures. IC chip needs will grow rapidly after the initial startup. Planned throughput will be more than adequate to meet near-term needs, so the facility could provide foundry services to support other SCE projects.

On-going advanced development will be required in areas such as increasing J_c , reducing feature sizes, planarization, etc.

TABLE 4-10. SUPERCONDUCTING IC CHIP MANUFACTURING COSTS						
	2006	2007	2008	2009	2010	Total
Pilot Line (\$M)	34.2	6.2	6.5	6.5	6.5	59.9
Manufacturing Line (\$M)	1.3	7.0	5.5	24.4	7.1	45.0
Support Activities (\$M)	5.1	3.9	2.4	6.1	2.7	20.2
Total Funding (\$M)	40.5	17.1	14.3	36.9	16.3	125.1
Total Staffing (FTE)	25	32	36	40	41	174

Table 4-10 summarizes the cost of IC chip manufacture. The first two lines summarize the year-by-year costs associated with the pilot and manufacturing lines. These include tooling purchases (with installation), operating expenses, and personnel. These costs are based on modifying the NGST foundry process as it existed prior to shut-down in 2004, but they should be representative regardless of where the new facility is placed or who operates it. Cost estimates are based on rounded 1999 quotations from tool vendors. The third row of the table summarizes the cost of support activities, such as packaging and testing, including personnel.

In years 1 and 4, the project incurs large expenses due to equipment purchases and facility upgrades for pilot and manufacturing lines, respectively. In year 5, additional process tools are brought online to increase throughput, if needed. The non-recurring costs associated with the pilot and manufacturing lines are \$59.9 million (M) and \$45M, respectively, over the life of the program, assuming that all process tools are purchased new. The manufacturing line is cheaper, because it builds on the existing pilot line.

A potential savings of ~40% in the cost of the fabrication tools could be realized if factory refurbished fabrication tools are available. Unfortunately, refurbished tools will not necessarily be available when needed by the project; therefore some combination of new and refurbished tools is the most likely scenario. (These costs assume use of existing facilities.) Upgrade to Class 1 would be additional expense, should it be warranted. Existing clean room facilities, unless built in the last few years, were designed for ballroom-style tools. To accommodate the new bulkhead-style tools may require structural modifications that could increase costs significantly. Fortunately, the wafer throughput requirements are low compared to commercial CMOS fabrication facilities, and relatively little duplication of tools is required, which minimizes the size of the clean room area. The pilot and manufacturing line cost figures are for fabrication only and do not include design and testing.

The nonrecurring design and test costs of complex superconductive IC chips are not well defined at present.

The IC chip manufacturing facility will also need to be able to package functional chips for further use for micro-architecture development, CAD development, and circuit design development, as well as multi-chip demonstrations. NGST had developed a carrier/MCM production and packaging process suitable for packaging a handful (<10) IC chips for multi-gigabit testing. A modest increase in layers and additional equipment for assuring reliability would probably be sufficient to support a 2010 machine. Table 4-10 includes estimates for the costs of testing, data extraction and reduction, and packaging for the IC chip manufacturing facility.

The nonrecurring design and test costs of complex superconductive IC chips are not well defined at present. The largest unknown cost is wafer level testing at 4 K, which will require implementation of cryogenic wafer probe technology or use of multi-chip testers to increase throughput. Cryogenic probe stations are estimated to cost about \$1M each and would perform the parametric testing and low speed functional circuit testing to prescreen the IC chips. At 10% yield, a throughput of 16 wafers per day is required. The number of probe stations needed depends on many factors but mainly on the time required to reach 4 K. In addition to the cryo-mechanics, each functional test station will require a full complement of instrumentation. An optimum functional IC chip test and hardware strategy needs to be developed to keep costs under control. An overall estimate can be developed using our estimates for costs and throughput of cryo wafer probers and the experience of NGST, which had a group dedicated to parametric testing and data extraction.

At the current level of spending on superconductive digital IC chip manufacturing, it is certain that chips of the density required for petaflops-scale digital systems will not become available in the foreseeable future.

The panel has described an aggressive program to achieve yield, density, and performance of IC chips suitable for petaflops-scale computing. As indicated in Table 4-10, the total estimated cost associated with IC chip manufacture over a five-year program is \$125.1M. A significantly reduced level of funding of roughly half of this figure could enable the establishment only of pilot line capabilities with manufacturing yield low or unknown, and less-than-desired IC chip densities. Clock rates would probably be increased to targets discussed in this study, as this is an area of high research interest. However, the processes used to achieve the devices for 50 GHz and greater clock rates would most likely use low-volume steps (e.g., e-beam lithography) and not be amenable to scaling for production.

It is estimated that current funding for superconductive-digital-IC chip manufacture, even at an R&D level, is far less than \$2M per year when both industrial and government funding are taken into account. At the current level of spending on superconductive-digital-IC chip manufacturing, it is certain that chips of the density required for petaflops-scale digital systems will not become available in the foreseeable future.

05

Packaging and chip-to-chip interconnect technology should be reasonably in hand.

Wideband data communication from low to room temperature is a challenge that must be addressed.

Total investment over five-year period: \$92 million.

INTERCONNECTS AND SYSTEM INPUT/OUTPUT

In petaflops-scale computer systems, the processor to memory and processor to processor data rates are enormous by any measure. The hybrid technology multi-thread (HTMT) program estimated the bisectonal bandwidth requirement to be 32 petabits/s. This bandwidth is composed of communications from relatively distant secondary storage, low-latency communications with “nearby” primary storage, and communications between processor elements within the cryogenic Rapid Single Flux Quantum (RSFQ) processing units.

The use of cryogenic RSFQ digital circuits with clock frequencies exceeding 50 GHz imposes challenges resulting from the increasing differential between memory cycle time and processor clock. Reduced time-of-flight (TOF) latency motivates the use of cryogenic memory close to the processor. Providing the required bandwidth between room-temperature electronics and the cryogenic RSFQ processor elements requires careful engineering of the balance between the thermal load on the cryogenics and the number, type, bandwidth, and active elements of the lines providing input/output (I/O).

The major interconnection, data communication, and I/O needs of a petaflops-scale system based on cryogenic RSFQ technology are:

- High throughput data input to the cryogenic processors and/or memory at 4 K.
- High throughput output from the 4 K operating regime to room-temperature system elements such as secondary storage.
- Communication between processor elements within the 4 K processing system at data rates commensurate with the processor clock rate.

In order to minimize the line count (thereby minimizing thermal and assembly issues) yet provide the requisite capability to carry tens of petabits/s of data, the bandwidth of each individual data line should be at least tens of Gbps.

Optical technologies offer a potential solution to this requirement while offering much lower thermal loads than Radio Frequency (RF) electrical cabling. However, optical components need to overcome issues of power, cost, and cryogenic operation. The low voltage levels of superconductive electronic devices are a challenge to direct communication of RSFQ signals from 4 K to room temperature. Some combination of Josephson junction (JJ), semiconductor, and optical components operating at a variety of temperatures will probably be required to provide the combination of data rate, low line count, and low thermal load required.

In a hierarchical supercomputer, such as any RSFQ petaflops system is likely to be, processor performance can be improved by data communication and manipulation at different levels of the memory hierarchy. As an example, HTMT provided a great deal of data manipulation and communication at the room temperature level using an innovative interconnection architecture called a Data Vortex. The Lightwave Research Laboratory at Columbia University has recently demonstrated a fully implemented 12-port Data Vortex Optical Packet Switch network with 160Gbps port bandwidth yielding a nearly 2Tbit/sec capacity.

Another challenge is the low latency (ns range) requirement imposed for data movement within the thousands of processors and memories at cryogenic temperature. Superconducting switches offer a potential solution for direct interfacing between RSFQ based processors and their immediate memory access. Table 5-1 summarizes the technologies, issues, and development for the data communication requirements of a cryogenic RSFQ petaflops-scale system. The panel estimates that the roadmap for these elements will require \$91.5 million (M), with a large team working some parallel paths early, then focusing down to a smaller effort demonstrating capability required for petaflops-scale computing by 2010.

TABLE 5-1. I/O TECHNOLOGIES FOR PETAFLOPS-SCALE PROCESSING WITH RSFQ

Data Communication Requirement	Technology	Status	Projections/Needs
Room Temperature to Cryogenic (Input)	Direct Optical to RSFQ	<ul style="list-style-type: none"> – 40-50 Gbps components available off the shelf, few demonstrations at cryo. – 20 Gbps RT to RSFQ demonstrated in the 90's. 	<ul style="list-style-type: none"> – Best candidate. – Need to demonstrate and qualify a manufacturable link. – Need to confirm direct optical to RSFQ conversion.
	Optical with OE conversion at 40-77 K	<ul style="list-style-type: none"> – 40-50 Gbps components available off the shelf, no demonstrations at cryo. 	<ul style="list-style-type: none"> – Fallback candidate. – Power associated with cryo-OE conversion is an issue.
	Direct Electrical input	<ul style="list-style-type: none"> – Well understood from R&D and small-scale applications. – Cable technology restricts practical line rates to <20 Gb/s. 	<ul style="list-style-type: none"> – Lowest risk candidate, but line count may become prohibitive for petascale. – Need to develop low thermal loss, high-bandwidth cables.
Cryogenic to Room Temperature (Output)	Direct RSFQ to Optical	<ul style="list-style-type: none"> – Some very-low-data-rate experiments have been done. – Fiber from 4 K provides lowest thermal loads, but heat from optical components at 4 K is a concern. 	<ul style="list-style-type: none"> – Highest risk approach, but development of low-power, high-speed modulators, if successful, provides lowest power and line count. – Needs research/development of innovative low-power cryogenic optical components.
	RSFQ to intermediate temperature electrical to optical out	<ul style="list-style-type: none"> – Electrical amplifiers demonstrated with low (5 mw) power at ~10 Gb/s. 	<ul style="list-style-type: none"> – Modularity, assembly, cost per line out will be a major factor. – Needs development of moderate (10 mW) power optical and electrical components at intermediate (40-77 K) temperatures.
	RSFQ Electrical Out	<ul style="list-style-type: none"> – Commonly used for 10 Gb/s in R&D and small-scale applications (coax). – Trades for custom cables and intermediate temperature amplifiers well understood. 	<ul style="list-style-type: none"> – Lowest risk approach, but line count may be prohibitive. – Needs aggressive development of RSFQ amplifiers/cryogenic semiconductor amplifiers and cabling to optimize power and bandwidth per line.
Cryogenic to Cryogenic	Superconducting switch	<ul style="list-style-type: none"> – 16x16 crossbar switch chips demonstrated at 20 Gb/s. – 4x4 Banyan elements demonstrated at 40 Gb/s. 	<ul style="list-style-type: none"> – Contention, latency, and architecture suitable for petascale need to be demonstrated.

Integrated, word-wide optical transmitter and receiver arrays are needed for all optical links.

5.1 OPTICAL INTERCONNECT TECHNOLOGY

This section describes the general properties and status of optical interconnect technology as it applies to all of the areas shown in Table 5-1. Discussions of the specific developments required in each of these areas are in the sections that follow.

While RSFQ processors allow construction of a compact ($\sim 1 \text{ m}^3$) processing unit, a superconducting petaflops-scale computer is a very large machine, on the scale of tens of meters, with high data bandwidth requirements between the various elements. For example, a particular architecture may require half a million data streams at 50 Gbps each between the superconducting processors and room-temperature SRAM. One potential solution is the use of optical interconnect technologies.

*For a superconductive supercomputer,
the very low thermal conductivity of glass fibers
substantially reduces the wall-plug power
of the cooler.*

The major advantage of optics is in the regime of backplane or inter-cabinet interconnects where the low attenuation, high bandwidth, and small form factor of optical fibers become valuable. The use of optical fibers using Wavelength Division Multiplexing (WDM) reduces the large number of interconnections. A comparison between electrical and optical transmission indicates that for data rates higher than 6-8 Gbps, the distance that electrical transmission is advantageous over optical interconnects does not exceed 1 meter. For a superconductive supercomputer, the thermal advantages of glass vs. copper are also very important. The very low thermal conductivity of the 0.005" diameter glass fibers, compared to that of copper cables necessary to carry the same bandwidth, represents a major reduction on the heat load at 4 K, thereby reducing the wall-plug power of the cooler substantially.

The term "optical interconnect" generally refers to short reach ($< 600 \text{ m}$) optical links in many parallel optical channels. Due to the short distances involved, multimode optical fibers or optical waveguides are commonly used. Optical interconnects are commercially available today in module form for link lengths up to 600 m and data rates per channel of 2.5 Gbps with a clear roadmap to 10-40 Gbps. These modules mount directly to a printed circuit board to make electrical connection to the integrated circuits, and use multimode optical ribbon fiber to make optical connection from a transmitter module to a receiver module. Figure 5.1 illustrates how optical interconnects might be used for both input and output between 4 K RSFQ and room-temperature electronics.

For further discussions of the options considered, see Appendix K: Data Signal Transmission. (The full text of this appendix can be found on the CD accompanying this report.)

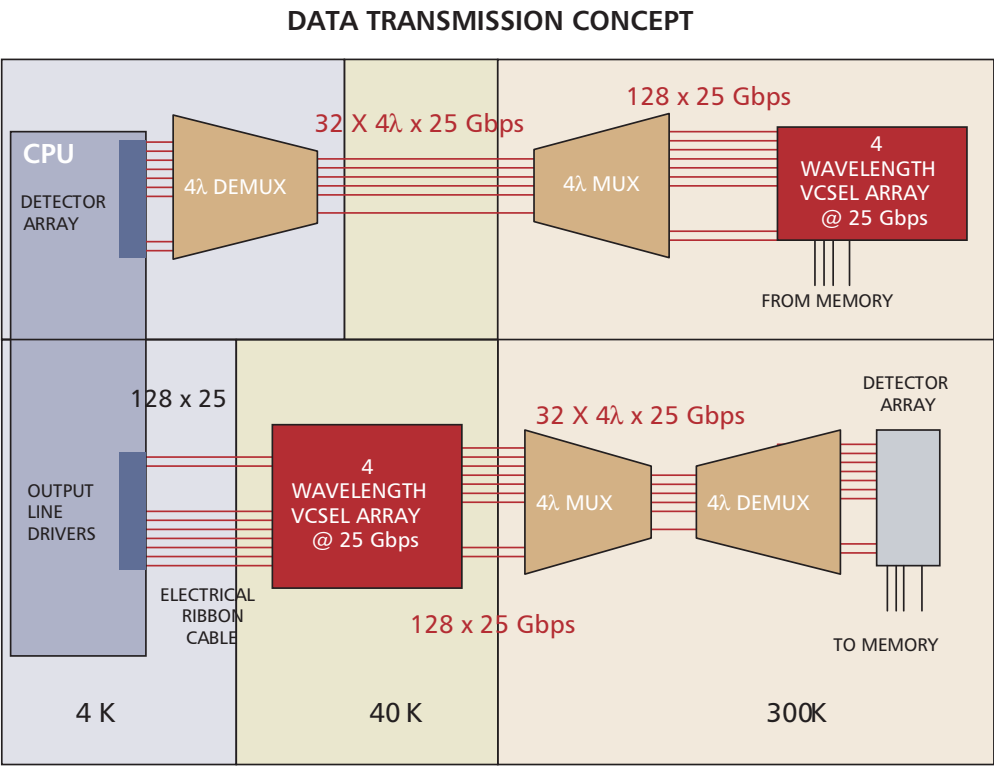


Figure 5-1. A 64-fiber, 4-wavelength, 25-Gbps CWDM System for bi-directional transmission totaling 6.4 Tbps between a superconducting processor at 4 K and high speed mass memory at 300 K. Optical connections are shown in red, electrical in black. This technology should be commercially available for 300 K operation by 2010.

5.1.1 OPTICAL INTERCONNECT TECHNOLOGY – STATUS

The need to move optical interconnects closer to the I/O pin electronics requires advances in packaging, thermal management, and waveguide technology, all of which will reduce size and costs. The research is ongoing with some efforts funded by industry, and others by governmental entities such as Defense Advanced Research Projects Agency (DARPA). For example, using Vertical Cavity Surface Emitting Lasers (VCSELs) and Coarse WDM, the joint IBM/Agilent effort has achieved 240 Gbps aggregate over 12 fibers, each carrying four wavelengths at 5 Gbps each (Figure 5-1). The next step is planned to be 480 Gbps, with each channel at 10 Gbps. Optical interconnects were already used in some large scale commercial high-performance data routers to connect backplanes and line cards together.

TABLE 5-2. ACADEMIC ACTIVITY

Univ. of Colorado	Plan to build an array of 40 Gbps directly modulated VCSELs; has already achieved 10 Gbps and expect to reach 20 Gbps in the near future.
Univ. of Texas at Austin	Working on amplifierless receivers crucial for 4 K applications.
Univ. of California at Santa Barbara	Working on amplifierless receivers crucial for 4 K applications.
All of these efforts are aimed at achieving devices and systems that will be practical for use in the near future in terms of cost and power consumption as well as performance.	

5.1.2 OPTICAL INTERCONNECT TECHNOLOGY – READINESS

Although optical interconnects have been introduced in large-scale systems, some technical issues must be resolved for the use of optical interconnects in a superconducting computer:

- The cryogenic operation of optical components must be determined; all the current development efforts are aimed at room temperature use with only minimal R&D effort for use at cryogenic temperatures. There do not appear to be any fundamental issues which would preclude the necessary developments.
- Even with cryogenic operation, some optical components may dissipate too much power at the required data rates to be useful in a petaflops-scale system. However, Figure 5.2 shows that low power is achievable at significant data rates for many optical components.

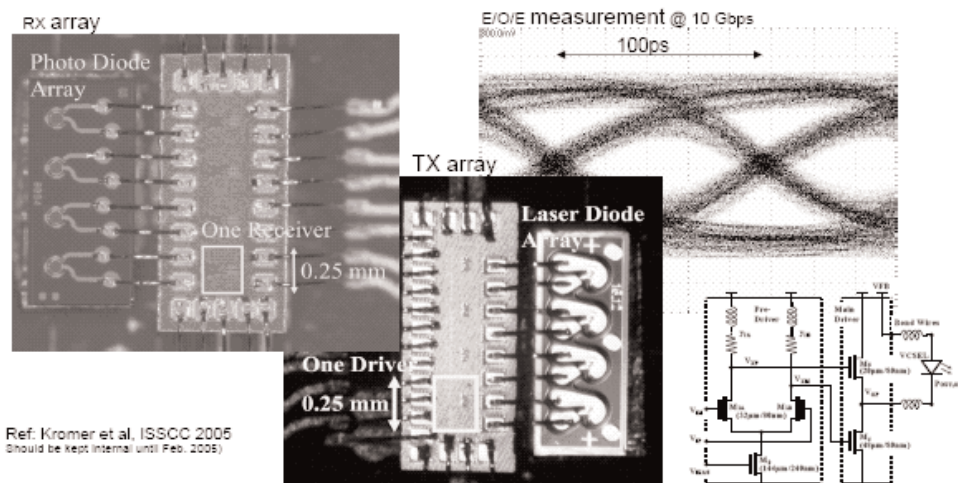


Figure 5-2. Four channel transceiver arrays operating at 10 Gbps/channel produced by IBM/Agilent. Total power consumption for 4 channels is less than 3 mW/Gbps.

5.1.3 OPTICAL INTERCONNECT TECHNOLOGY – PROJECTIONS

The availability of optical components and their features is given in Table 5-3.

TABLE 5-3. PROJECTIONS FOR ELECTRO-OPTICAL COMPONENTS FOR RSFQ I/O				
DEVICE	CURRENT STATUS 2005	DESIRED STATE 2010	CRYOGENIC USE TESTED	ONGOING R&D
VCSEL Based CWDM Interconnect Systems: 12x4λ @ 300K				
DARPA Terabus	12 x 4λ x 10 Gbps	8 x 8λ x 50 Gbps	N	Y
LIGHT SOURCES				
50 Gbps VCSEL Arrays	10 Gbps	50 Gbps	Y	Y
Frequency Comb Lasers	32λ x 1550 nm	64λ x 980 nm	N/A	N
OPTICAL MODULATORS				
VCSEL used as DWDM modulator	2.5 Gbps	50 Gbps	N	Y
Low Voltage Amplitude Modulator @ 40 K	4-6 volts On-Off/ 40 Gbps Single Devices/ COTS	0.5 volt on-off/ 50 Gbps Arrays	N	Y
Exotic Modulator Materials: Organic Glass/Polymer/Magneto-Optic	Materials studies show 3x improvement over conventional materials	6-10 fold improvement in devices	N	Y
OPTICAL RECEIVERS				
Amplifierless Receivers matched to superconductors	20 Gbps demonstrated	50 Gbps Array	N	N
25/50 Gbps Receiver Arrays @ 300K	50 Gbps Single Devices – COTS 4 x 40 Gbps arrays Lab	Ideally, word-wide arrays	N/A	N
ELECTRONIC DEVICES				
Low Power, 50 Gbps Amplifier/Driver Arrays @ 40 K	Single Amplifiers (InP-HEMT) – COTS	Lower power, ideally, word-wide arrays	Y	N
PASSIVE COMPONENTS				
Cryogenic Electrical Ribbon Cables at 25/50 Gbps	3 Gbps	25/50 Gbps	Y	N
Optical Ribbon Cables @ 4 K, 40 K	COTS meet Commercial Specs/ Fibers OK at 4 K	Cables at 4 K	N	N
Optical WDM Components @ 4 K, 40 K	COTS meet Commercial Specs	4 K Operation	N	N

5.1.4 OPTICAL INTERCONNECT TECHNOLOGY – ISSUES AND CONCERNS

Commercially driven telecommunications technology is already producing optical and electronic components capable of operating at 40 Gbps data rates, although these components are still very expensive, considering the 500,000+ individual channels required for a petaflops-scale system. (A single 40 Gbps transmitter line card—electrical data in/optical out—costs between \$10,000 and \$20,000.) Also, telecommunications economics demand that these transmitters be far more powerful than is needed for this application. The telecommunications market will not meet RSFQ needs to:

- Drive the cost down several orders of magnitude.
- Settle for lower power (both optical and electrical) solutions such as VCSELs.
- Provide word-wide integrated packages.

To have the necessary optical components available by 2010 will require investments to extend the results of the DARPA Terabus Program to the 25-50 Gbps/channel requirement as well as address the issues raised by cryogenic operation. The key issues to be addressed are:

- Developing device structures capable of 25-50 Gbps operation in highly integrable formats.
- Reducing electrical power requirements of large numbers of electro-optical components operating at cryogenic temperatures.
- Packaging and integration of these components to reduce the size and ease the assembly of a system of the scale of a petaflops-scale supercomputer.
- Reducing the cost of these components by large-scale integration and simple and rapid manufacturing techniques.
- Evaluating and adapting both passive and active commercially available optical components for use at cryogenic temperatures.

For a detailed roadmap of the optical developments required, see Appendix K: Data Signal Transmission. (The full text of this appendix can be found on the CD accompanying this report.)

The following sections detail the technology developments needed for each of the specific Data Communications requirements outlined in Table 5-1.

5.2 INPUT: DATA AND SIGNAL TRANSMISSION FROM ROOM TEMPERATURE TO 4 K

Direct optical transmission from 300 K to 4 K can be accomplished in a variety of ways and requires the least development beyond that already planned under DARPA programs. The major issues here will be the:

- Need to produce low cost, manufacturable transmitter components operating ideally at the 50 Gbps processor clock speed.
- Requirement to certify optical WDM components for cryogenic operation.

In this case, an optical receiver would be placed at 4 K to capture the data flow into the processor. Reasonable detected optical powers ($200\mu\text{W}$) would generate adequate voltage (100-160 mV) to drive the superconducting circuitry directly if a 100 ohm load and appropriate thresholding can be done. Alternately, since superconductive circuits are current driven, the superconducting logic could be driven directly with the detected current (100-160 μA) from the photodiodes. If this is achievable, it may be very advantageous to have the 300 K to 4 K interconnect be all optical, even if the reverse path is not. However, if any electronic amplification were required, the amplifier power would likely become a major thermal problem. This issue must be closely examined.

5.2.1 INPUT: ROOM TEMPERATURE TO 4 K – STATUS

Direct optical transmission from room temperature to 4 K, using a simple MSM diode to convert the input photons to the RSFQ level signals (0.1 to 0.2 mA) required has been shown in R&D demonstrations at HYPRES which have achieved 10's of Gbps serial data rates. It remains to validate increased data rates and integration with an RSFQ VLSI chip.

5.2.2 INPUT: ROOM TEMPERATURE TO 4 K – ISSUES AND CONCERNS

An issue which must be thoroughly explored is the use of room-temperature optical components in a cryogenic environment. Optical fiber has routinely been used at low temperatures in laboratory and space environments. However, telecommunications-grade optical components, such as fiber ribbon cables, connectors, wave division multiplexers, etc., must be evaluated for use at low temperatures.

The initial phases of a funded development effort should focus on identifying those areas which require adaptation of standard room-temperature optical and opto-electronic technologies for use at low temperatures. This is included in Table 5-4.

5.2.3 INPUT: ROOM TEMPERATURE TO 4 K – ROADMAP

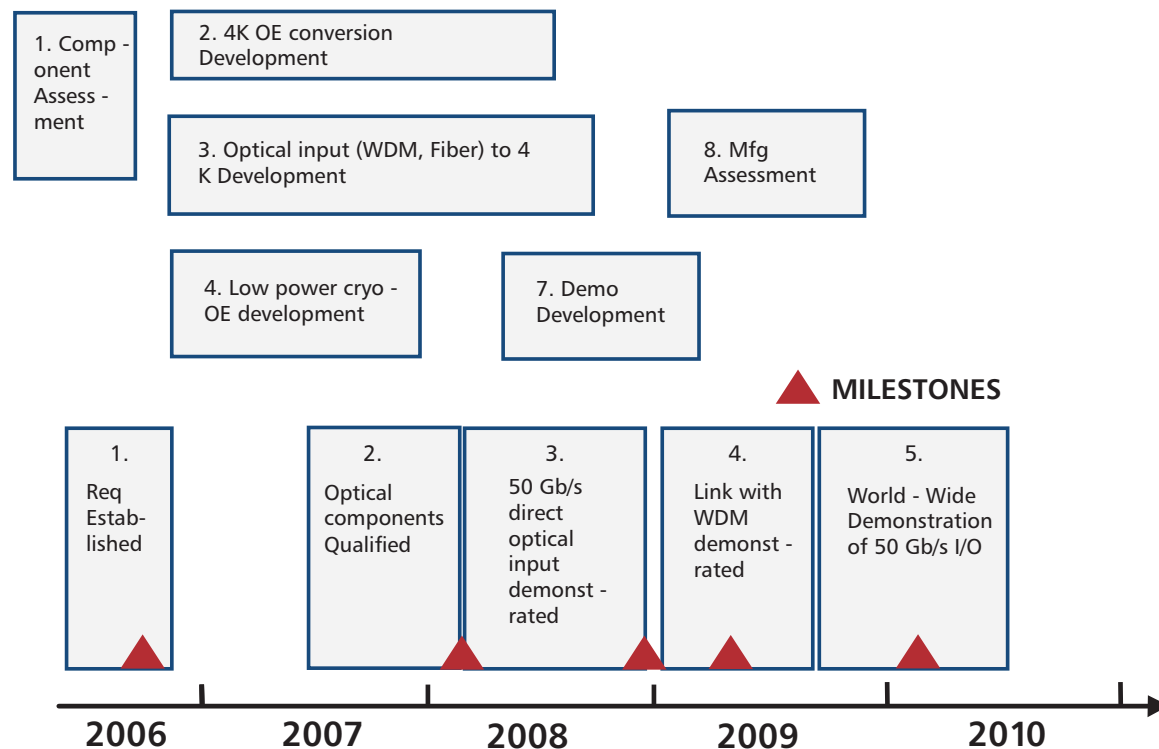


TABLE 5-4. FUNDING FOR ROOM TEMPERATURE TO 4 K ELECTRONICS

Element	2006	2007	2008	2009	2010	Total
Input (RT to 4 K)						
Manpower (FTE)	14	10	8	4	4	40
Funding (\$M)	6.0	3.8	2.8	3.0	1.1	16.7

5.3 OUTPUT: 4 K RSFQ TO ROOM TEMPERATURE ELECTRONICS

Output interfaces are one of the most difficult challenges for superconductive electronics. The low power consumption of RSFQ is a virtue for high-speed logic operation, but it becomes a vice for data output: There is not enough signal power in an SFQ data bit to communicate directly to conventional semiconductor electronics or optical components. Interface circuits are required to convert a voltage pulse into a continuous voltage signal of significant power. Josephson output circuits for electrical data transmission have been demonstrated at data rates up to 10 Gbps, and there is a reasonable path forward to 50 Gbps output interfaces. However, the balance of JJ, semiconductor, and optical components and their location (in temperature) is a critical design issue.

Given the very low power consumption of superconducting processors, the total thermal load from the inter-level communications becomes a major issue in the system. Given the power efficiency of refrigeration systems, placing transmitting optics at 4 K does not appear to be feasible.

Another option is to generate laser light at room temperature and transmit it by fiber into the 4 K environment where it would be modulated by the data on each line and then transmitted on to the next level. This would be a very attractive option if an optical modulator capable of operating at 50 Gbps, with a drive voltage (3-5 mV) consistent with superconducting technology existed. Currently available devices at 40 GHz require drive voltages in the range of 6 V_{p-p}, with rf drive powers of 90 mW. Furthermore, the device properties and the transmission line impedances available are not consistent with very low power consumption.

Since placing optical output components at 4 K is unattractive, other approaches must be considered. The design of the cryogenics appears to be compatible with having a short (<5 cm) section of high speed flexible ribbon cable, connecting the 4 K section of the system with a section at some intermediate temperature. The issue of the incompatibility of voltage levels between superconducting circuitry and optics can now be addressed using advanced technology amplifiers, such as InP HEMT transistors, which have been shown to operate at low temperatures with adequate bandwidth.

5.3.1 OUTPUT: 4 K RSFQ TO ROOM TEMPERATURE ELECTRONICS – STATUS

The demonstrations of data output from Josephson logic to semiconductor electronics have employed four basic techniques:

- **Stacked SQUIDs.** SFQ/DC converters develop only 150 μ V each, but series stacks develop a millivolt.
- **SFQ/Latch converters.** The SFQ pulse triggers latching junctions to develop 2 mV.
- **Suzuki stacks.** One latched junction triggers series arrays to develop 10 mV.
- **Suzuki stacks with GaAs amplifiers** on a multi-chip module (MCM) at 4 K.

Examples of what has been achieved through each approach are given in Table 5-5 below. Josephson output interfaces have demonstrated Gbps communication of data to room-temperature electronics. Actual bit error rates are better than listed for cases where no errors were found in the data record. Stacked SQUIDs use DC power and produce Non-Return to Zero (NRZ) outputs, which are significant advantages, but at the cost of many more JJs per output. Latching interfaces develop higher voltages than SQUIDs, but require double the signal bandwidth for their RZ output. The latching interface must also be synchronized to the on-chip data rate.

Higher performance is expected for higher-current-density junctions; speed increases linearly for latching circuits and increases as the square root of current density for non-latching circuits. The Advanced Technology Program demonstration at NGST (then TRW) successfully integrated semiconductor amplifiers (10 mW dissipation) onto the same multi-chip module with Suzuki stack outputs.

Both coaxial cable and ribbon cable have been used to carry signals from 4 K to room-temperature electronics. In the NSA crossbar program, GaAs HBT amplifiers (10-30 mW dissipation) were operated at an intermediate temperature, approximately 30 K.

TABLE 5-5. RESULTS OF OUTPUT TECHNIQUES

JJ output type	JJ count	dc power	V _{out} (mV)	NRZ	Rate (Gbps)	BER (max)	J _c (kA/cm ²)
Stacked SQUIDs	60	Yes	1.3	Yes	1.0	1e-07	1
SFQ/Latch	5	No	2	No	3.3	1e-08	8
Suzuki Stack (6X)	17	No	12	No	10	1e-07	8
Suzuki (4X) + GaAs Amplifier	12	No	10	No	2	1e-09	2

5.3.2 OUTPUT: 4 K RSFQ TO ROOM TEMPERATURE ELECTRONICS – READINESS AND PROJECTIONS

Commercial off-the-shelf (COTS) fiber optic components provide much of the basic infrastructure for input and output between room-temperature and 4 K electronics. For example, 40 Gbps (OC-768) transceiver parts are becoming widely available. However, as discussed in Section 5.1, they are very expensive and are not designed to be compatible with word-wide, low-power, short-range applications. A significant effort must be put into tailoring COTS systems to meet the requirements of an RSFQ computer.

The challenge for output interfaces is raising the signal level by 60 dB from 200 μ V to 200 mV at 50 Gbps. HEMT amplifiers have demonstrated low power (4 mW) operation at 12 K with low noise figures of 1.8 dB in the band 4-8 GHz. Modern High Electron Mobility Transistor amplifiers are capable of the 35 GHz bandwidth needed for NRZ outputs at 50 Gbps, and will operate at 4 K.

A significant effort must be put into tailoring COTS systems to meet the requirements of an RSFQ computer.

Josephson output interfaces have demonstrated low bit error rate at 10 Gbps, using 8 kA/cm² junctions. A critical current density of 20 kA/cm² is probably sufficient for 50 Gbps output interfaces.

The output voltage of demonstrated JJ interface circuits is sufficient to sustain a low bit error rate at 50 Gbps. Reasonable estimates can be made of the signal power required at the Josephson output drivers. If ribbon cable can carry electrical signals with 3 dB of signal loss to amplifiers at 40 K with 3 dB noise figure, then 4 mV_{pp} on the Josephson drivers will sustain a bit error rate of 1e-15.

The crossbar switch program provides a rich experience base in the art of ribbon cables. In particular the trade off between heat leak and signal loss is well understood.

5.3.3 OUTPUT: 4 K RSFQ TO ROOM TEMPERATURE ELECTRONICS – ISSUES AND CONCERNS

A focused program to provide the capability to output 50 Gbps from cold to warm electronics must overcome some remaining technical challenges, both electronic and optical.

Electronics Issues

There are two major electronic challenges:

- Designing custom integrated circuits in a high-speed semiconductor technology to minimize refrigeration heat loads. These include analog equalizers to compensate for frequency dependent cable losses, wideband low-noise amplifiers, and 50 Gbps decision circuits. COTS parts are optimized for room-temperature operation and they run very hot. ASICs optimized for low power operation at cryogenic temperatures will be needed. The circuitry which directly drives electro-optical components such as VCSELs or modulators must be easily integrable with these device technologies.
- Designing ribbon cables with better dielectrics to carry signals at 50 Gbps. Existing cables have significant dielectric losses at frequencies above 10 GHz. Data at 50 Gbps has significant power at frequencies up to 35 GHz for NRZ format, and up to 70 GHz for RZ format. The development of better RF cables, particularly for 4 K to intermediate temperature, is discussed in Section 6.

Optics Issues

Although development of components to either generate or modulate optical streams at 50 Gbps at cryogenic temperatures will require substantial investment, the panel believes current efforts in the field support optimism that one or more of the possible approaches will be successful by 2010. In decreasing order of difficulty these efforts are:

- Achieving low drive power, word-wide arrays of VCSELs, capable of operating at 50 Gbps either as a directly modulated laser or as an injection-locked modulator.
- Achieving 50 Gbps word-wide receiver arrays complete with data processing electronics, for room-temperature operation.
- Producing a frequency comb laser to be used in conjunction with low-power modulators to allow use of Dense Wavelength Division Multiplexing to reduce optical fiber count to one input and one output fiber per processor, each carrying up to 4 Tbps.

5.3.4 OUTPUT: 4 K RSFQ TO ROOM TEMPERATURE ELECTRONICS – ROADMAP AND FUNDING PROFILE

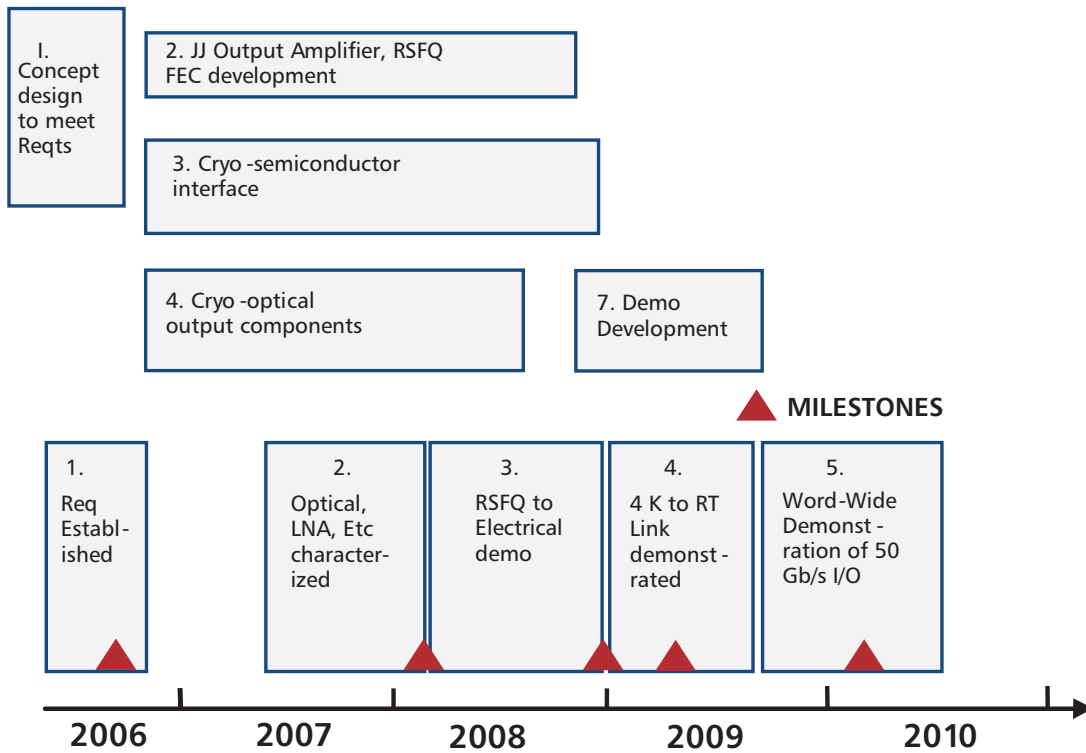


TABLE 5-6. FUNDING FOR 4 K TO ROOM TEMPERATURE ELECTRONICS

Element	2006	2007	2008	2009	2010	Total
Output (4 K to RT)						
Manpower (FTE)	23	30	27	24	13	117
Funding (\$M)	12.9	16.1	16.2	14.9	7.4	67.5

5.4 DATA ROUTING: 4 K RSFQ TO 4 K RSFQ

The interconnection network at the core of a supercomputer is a high-bandwidth, low-latency switching fabric with thousands or even tens of thousands of ports to accommodate processors, caches, memory elements, and storage devices. Low message latency and fast switching time (both in the range of ns), along with very high throughput under load and very-high line data rates (exceeding 50 Gbps), are the key requirements of the core switching fabric.

The crossbar switch architecture (developed on a previous NSA program, with low fanout requirements and replication of simple cells) can be implemented with superconductive electronics. A superconducting crossbar supplies a hardwired, low-latency solution for switches with high traffic and a large number of ports. The low row and column resistances allow scalability to hundreds or thousands of ports, and the use of high speed superconductive ICs permits serial data transmission at 50-100 Gbps, bypassing I/O-bound parallelization of serial data streams.

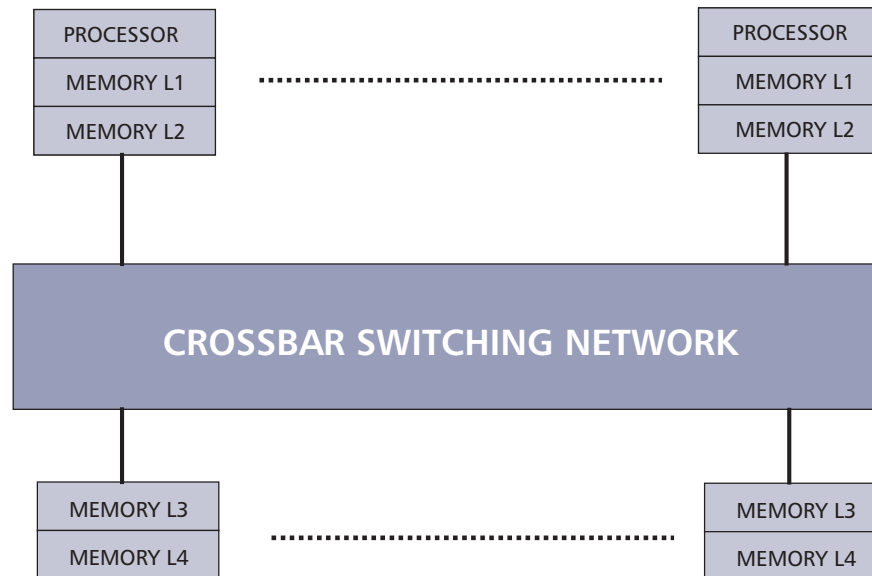


Figure 5-3. The use of crossbar switching in supercomputers.

5.4.1 DATA ROUTING: 4 K RSFQ TO 4 K RSFQ – STATUS

Electronic technology continues to make huge strides in realizing high-speed signaling and switching. Data rates of 10 Gbps have been realized on electronic serial transmission lines for distances in the order of tens of centimeters on printed circuits boards and as long as 5 meters on shielded differential cables. Application-specific integrated circuits (ASIC) with increasing speed and decreasing die area, used as switching fabrics, memory elements, and arbitration and scheduling managers, have shown increasing performance over a long period of time.

Issues such as power, chip count, pin count, packaging density, and transmission line limitations present a growing challenge to the design of high-rate semiconductor-based electronic switches. Next-generation semiconductor processes, along with increased device parallelism and implementation of multistage architectures, can increase the total switching capacity of a multistage semiconductor-based electronic switch to a 128 x 128 matrix of 50 Gbps ports. Progress beyond such switches is expected to be slow, as evidenced in the recent International Technology Roadmap for Semiconductors (ITRS) report.

Alternatively, superconductivity is an ideal technology for large-scale switches. Ideal transmission lines, compact assemblies, and extremely low gate latency permit operation with low data skew (hence permitting low overhead transmission of word-wide parallel data for extremely high throughput). The superconductive implementation of the crossbar architecture permits large switches, operating in single-cast or broadcast mode.

Components of a 128x128 superconductive scalable crossbar switch have been demonstrated at NGST (Figure 5-4), using voltage state superconductive switch elements. The switch chips were inserted onto an MCM along with amplifier chips, and data transmission up to 4.6 Gbps was observed. Latencies were in the order of 2-5 ns. Simulations indicated that this type of switch could be operated at >10 Gbps, but 50 Gb/s require RSFQ switch elements. The crucial component for viable RSFQ switches at 20-50 Gbps is cryogenic amplification without high power penalties.




Figure 5-4. Superconducting 16x16 crossbar switch with 14,000 junctions on a 5x5 mm chip (NGST).

Other superconductive switch architectures have been demonstrated, including Batcher-Banyan components and a prototype token-ring network element by NEC. NEC also demonstrated a small scale SFQ circuit for asynchronous arbitration of switch fabrics operating at 60 Gbps. The Batcher-Banyan architecture reduces the number of switch elements (growth as $N \log N$ instead of N^2), but the element interconnectivity requirements can make latency and skew management difficult. The token-ring architecture allows drastic reduction in the circuitry required for high-bandwidth switching but has inherent latency which must be overcome.

5.4.2 DATA ROUTING: 4 K RSFQ TO 4 K RSFQ – READINESS

The basic superconducting crossbar chips have been demonstrated and the scalability issues for the implementation of larger switching networks are well understood. But because higher data rates and lower latencies for petaflops-scale systems need more engineering, more development funding is needed to provide a large-scale, low-latency switching solution.

5.4.3 DATA ROUTING: 4 K RSFQ TO 4 K RSFQ – ISSUES AND CONCERNS

Although the realization of a fast switch using superconducting circuits is feasible, some issues remain for large-scale switches:

- **Memory:** Larger switching fabrics require on-board memory to help maintain the throughput of the switch under heavy traffic loads. The amount of memory is a function of the system and switch architecture. The lack of a well established superconducting memory is a concern for the scalability.
- **Processing logic:** Implementation of larger switch fabrics with minimal hardware cost requires on-board flow control (scheduling). The realization of scheduling circuits with superconducting circuits is not well established.
- **Line rates and port widths:** Depending on the system and switch fabric architectures the logical widths and data line rates at each switch port may exceed the limit of auxiliary I/O technologies especially on the side of the switch interfacing to the room temperature.

5.4.4 DATA ROUTING: 4 K RSFQ TO 4 K RSFQ – ROADMAP AND FUNDING

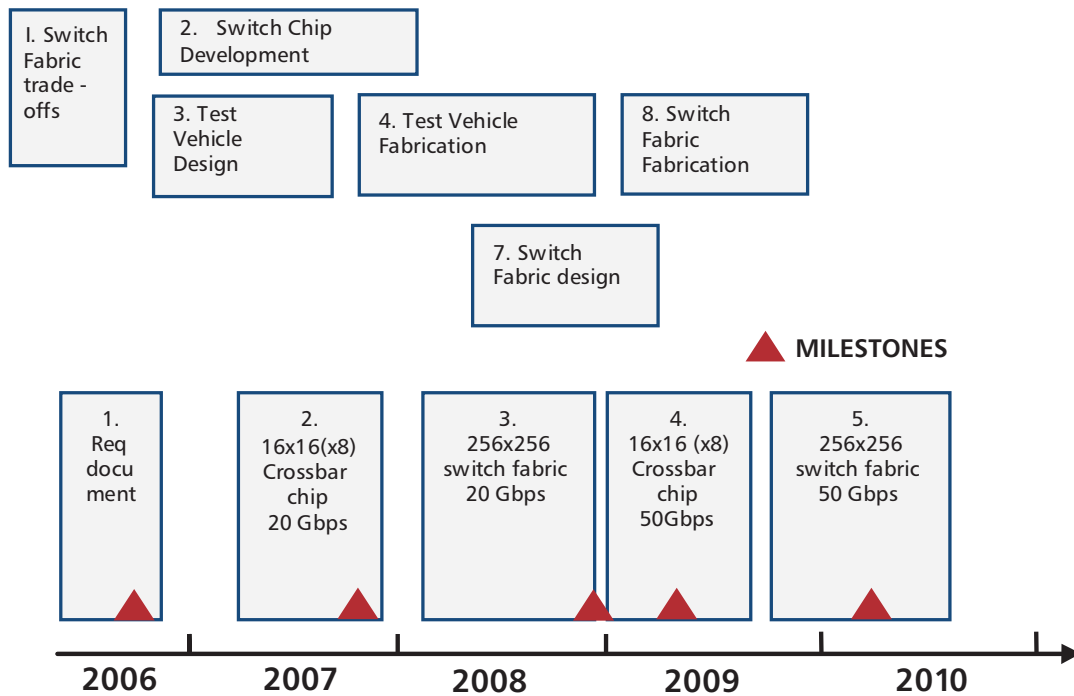


TABLE 5-7. 4 K RSFQ TO 4K RSFQ						
Element	2006	2007	2008	2009	2010	Total
Cryo-Data Routing (4 K)						
Manpower (FTE)	3	5	4	3	3	18
Funding (\$M)	0.9	1.4	3.0	1.0	1.0	7.3

06

The design of secondary packaging technologies and interconnects for SCE chips is technically feasible and fairly well understood.

The lack of a superconducting packaging foundry with matching design and fabrication capabilities is a major issue.

Enclosures, powering, and refrigeration are generally understood, but scale-up issues must be addressed.

System testing issues must be addressed.

Total investment over five-year period: \$81 million.


SYSTEM INTEGRATION

System integration is a critical but historically neglected part of the overall system design. Superconductive electronic (SCE) circuits offer several challenges due to their extremely high clock rates (50-100 GHz) and ability to operate at extremely cold temperatures. The ultra-low power dissipation of Rapid Single Flux Quantum (RSFQ) logic means that a compact package can be used, enabling high-computational density and interconnect bandwidth. The enclosure must also include magnetic and radiation shielding needed for reliable operation of SCE circuits.


System integration for large-scale SCE systems was investigated in previous programs such as HTMT¹, whose approach is shown in Figure 6.1. In this concept, the cryogenic portion of the system occupies about 1 m³ with a power load of 1 kW at 4 K. Chips are mounted on 512 multi-chip modules (MCM) that allow a chip-to-chip bandwidth of 32 Gbps per channel. Bisectonal bandwidth into and out of the cryostat is 32 Pbps.

The major components of this system concept are:

- Chip-level packaging including MCMs, 3-D stacks and boards.
- Cables and power distribution hardware.
- System-level packaging including enclosures and shields, refrigeration unit, system integration and test.



A)



B)

Figure 6.1. A) System installation concept for petaflops HTMT system. Enclosure for the superconducting processors is 1 m³ white structure with cooling lines into the top. B) Packaging concept for HTMT SCP, showing 512 fully integrated multi-chip modules (MCMs) connected to 160 octagonal printed circuit boards (PCBs). The MCMs, stacked four high in blocks of 16, are connected to each other vertically with the use of short cables, and to room temperature electronics with flexible ribbon cables (I/Os). The drawing has one set of the eight MCM stacks missing, and only shows one of the eight sets of I/O cables.

¹ HTMT Program Phase III Final Report, 2002

For this study, the readiness of the system integration technologies required by RSFQ was evaluated. The major conclusions are:

- The design of secondary packaging technologies (e.g., boards, MCMs, 3-D packaging) and interconnects (e.g., cables, connectors) for SCE chips is technically feasible and fairly well understood.
- The lack of a superconducting packaging foundry with matching design and fabrication capabilities is a major issue.
- The technology for the refrigeration plant needed to cool large systems, along with the associated mechanical and civil infrastructure, is understood well enough to allow technical and cost estimates to be made.
- Testing of a superconducting supercomputer has not been fully addressed yet. Testing mostly addressed modular approaches at cold temperatures by providing standard physical interfaces and limited functional testing at high frequencies.

6.1 MULTI-CHIP MODULES AND BOARDS

In order to accommodate thousands of processor and other support chips for a SCE-based, petaflops-scale supercomputer, a well-designed hierarchy of secondary packaging starting from RSFQ and memory chips and including MCMs (with very dense substrates supporting these chips) and printed circuit boards housing MCMs is needed. In addition, MCMs and boards are expected at an intermediate (40-77 K) for semiconductor electronics for data communications and potentially for memory. Figure 6.2 illustrates such a packaging concept developed for the HTMT design.

6.1.1 MULTI-CHIP MODULES AND BOARDS – STATUS

Demonstrations of RSFQ chip-to-chip data communication with low bit error rate (BER) at 60 Gbps were carried out encouraging the notion that a petaflops MCM technology with inter-chip communication at the clock rate of the chips is feasible (Figure 6.3). However, these demonstrations were carried out on very simple MCM-D packages with superconducting inter-chip connections, fabricated in an R&D environment.




Figure 6-2. HTMT conceptual packaging for cryogenic processing and data communications.




Figure 6.3. A multi-chip module with SCE chips(left: NGST 's Switch chip MCM with amplifier chip, center: NGST's MCM; right HYPRES' MCM).

6.1.2 MULTI-CHIP MODULES AND BOARDS – READINESS

The design of MCMs for SCE chips is technically feasible and fairly well understood. However, the design for higher speeds and interface issues needs further development. The panel expects that MCMs for processor elements of a petaflops-scale system will be much more complex, requiring many layers of impedance controlled wiring, with stringent crosstalk and ground-bounce requirements. While some of the MCM interconnection can be accomplished with copper or other normal metal layers, some of the layers will have to be superconducting in order to maintain low bit error rate (BER) at 50 GHz for the low voltage RSFQ signals. Kyocera has produced limited numbers of such MCMs for a crossbar switch prototype. These MCMs provide an example and base upon which to develop a suitable volume production capability for MCMs.

At intermediate temperatures, complex boards (some with embedded passive components) have been evaluated and tested for low-temperature operation. Although MCMs with SCE chips are feasible, the signal interface and data links impose further challenges. In particular, the large number of signal channels and high channel densities are difficult to achieve.

6.1.3 MULTI-CHIP MODULES AND BOARDS – PROJECTIONS

The technology projections for MCMs and PCBs are given in Table 6-1:

TABLE 6-1. TECHNOLOGY PROJECTIONS FOR MCMs AND PCBs				
Year	2001	2005	2007	2009
MCM				
Pad size (μm)	75	25	15	15
Pad pitch (μm)	125	75	30	30
Pad density (cm ⁻²)	2000	4000	6000	6000
Max. no. Nb layers	7	9	9	9
Max. no. W layers	30	40	40	40
Linewidth (μm)	3	3	2	2
Bandwidth (Gbps/wire)	20	30	40	40
Chip-to-chip SFQ	Yes	Yes	Yes	Yes
Chips per MCM	50	50	50-150	50-150
BACKPLANE				
Technology	Ceramic	Ceramic/ Flex	Flex	Flex
Size (cm ²)	30	50	100	100
Bandwidth (Gbps/wire)	5	10	25	50

6.1.4 MULTI-CHIP MODULES AND BOARDS – ISSUES AND CONCERNS

The panel expects that the technology will be available for such packaging, but the major issue to be addressed will be assuring a source of affordable packaging. Several approaches should be evaluated:

- Develop a superconducting MCM production capability, similar to the chip production capability (perhaps even sited with the chip facility to share facility and some staff costs). This is the most expensive approach, though perhaps the one that provides the most assured access.

- Find a vendor willing to customize its advanced MCM packaging process to include superconducting wire layers, and procure packages from the vendor. Because of the relatively low volumes in the production phase, the RSFQ development effort would have to provide most—if not all—of the NRE associated with this packaging. Smaller vendors would be more likely to support this than larger ones.
- Procure MCMs with advanced normal metal layers for the bulk of the MCM, then develop an internal process for adding superconducting wiring. This is less expensive than the first approach, but it depends on development of a process on vendor material, which may change without notices, so it affords us less assured access.

Additional important issues for manufacturability are:

- Minimizing parts count.
- System modularity (which allows individual chips or boards to be easily assembled and replaced).

6.1.5 MULTI-CHIP MODULES AND BOARDS – ROADMAP AND FUNDING

A roadmap for the development of MCMs with SCE chips is shown below. The funding profile needed to maintain development pace is listed in Table 6-2. The funding includes the establishment of a small superconducting MCM production capability. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)

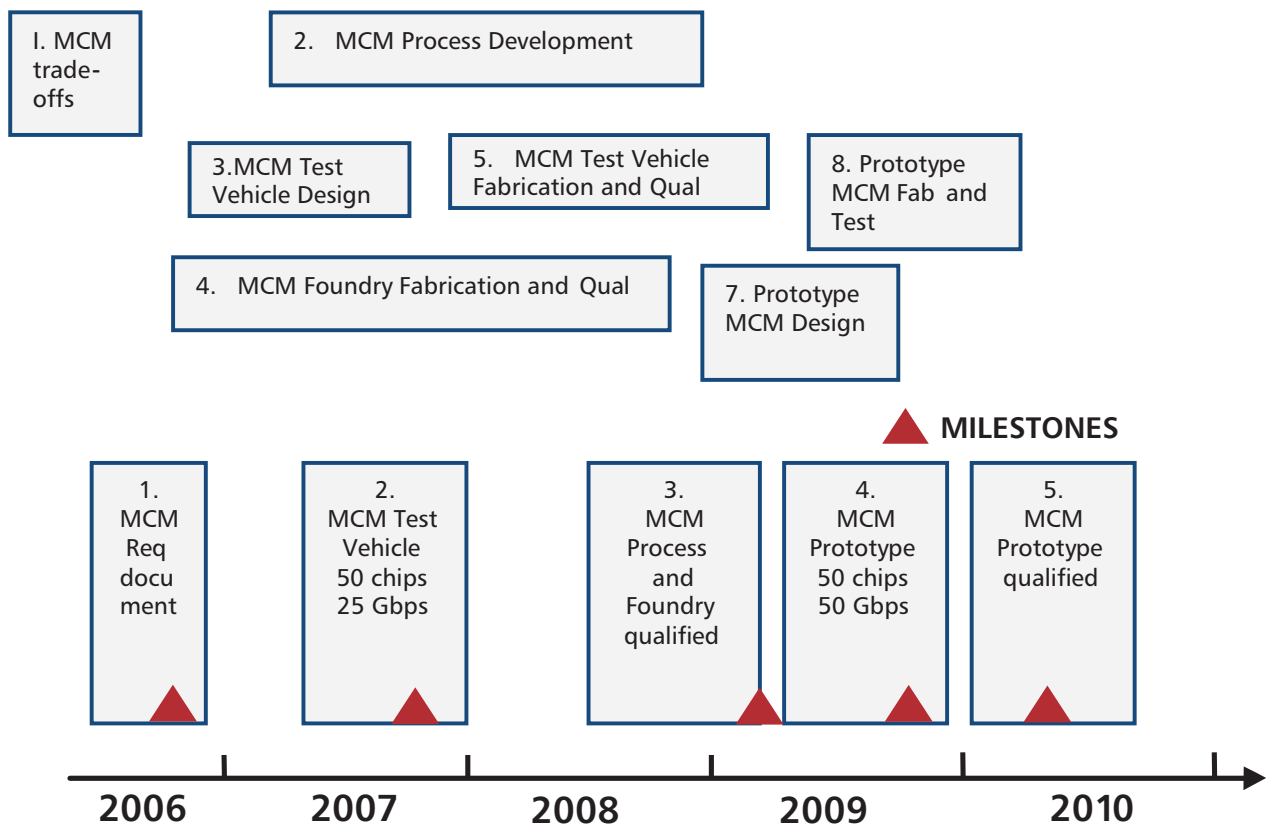


TABLE 6-2. MCM FABRICATION TOOLS AND DEVELOPMENT COSTS (\$M)

Year	2006	2007	2008	2009	2010	Total
Fab and Equipment Total	9.3	1.0	0.0	5.5	0.0	15.8
MCM Development Total	2.5	3.3	3.8	4.8	5.0	19.4
Total Investment	11.8	4.3	3.8	10.3	5.0	35.2

6.2. 3-D PACKAGING

Conventional electronic circuits are designed and fabricated with a planar, monolithic approach in mind with only one major active device layer along the z-axis. Any other active layers in the third dimension are placed far away from the layers below and there is no (or only a few) connection between layers. The trend of placing more active devices per unit volume resulted in technologies referred to as 3-D packaging, stacking, or 3-D integration. Compact packaging technologies can bring active devices closer to each other allowing short time of flight (TOF), a critical parameter needed in systems with higher clock speeds. In systems with superconducting components, 3-D packaging enables:

- Higher active component density per unit area.
- Smaller vacuum enclosures.
- Shorter distances between different sections of the system.

For example, 3-D packaging will allow packing terabytes to petabytes of secondary memory in few cubic feet (as opposed to several hundred cubic feet) and much closer to the processor.

6.2.1 3-D PACKAGING – STATUS

3-D packaging was initiated in the late 70's as an effort to improve the packaging densities, lower system weight and volumes, and improve electrical performance. Main application areas for 3-D packaging are systems where volume and mass are critical^{2,3}. Historically, focal-plane arrays with on-board processing and solid-state data recorder applications for military and commercial satellites have driven the development of 3-D packages for memories. Recently, 3-D packaging has appeared in portable equipment for size savings. 3-D packaging has expanded from stacking homogenous bare die (e.g., SRAM, Flash, DRAM) to stacking heterogeneous bare die and packaged die for "system-in-stack" approaches (Figure 6-4).

² S. Al-Sarawi, D. Abbott, P. Franzon, *IEEE Trans. Components, Packaging and Manufacturing Technology-Part B*, 21(1) (1998) p. 1

³ V. Ozguz, J. Carson, in *SPIE Optoelectronic Critical Review on Heterogeneous Integration*, edited by E. Towe, (SPIE Press, 2000), p. 225.

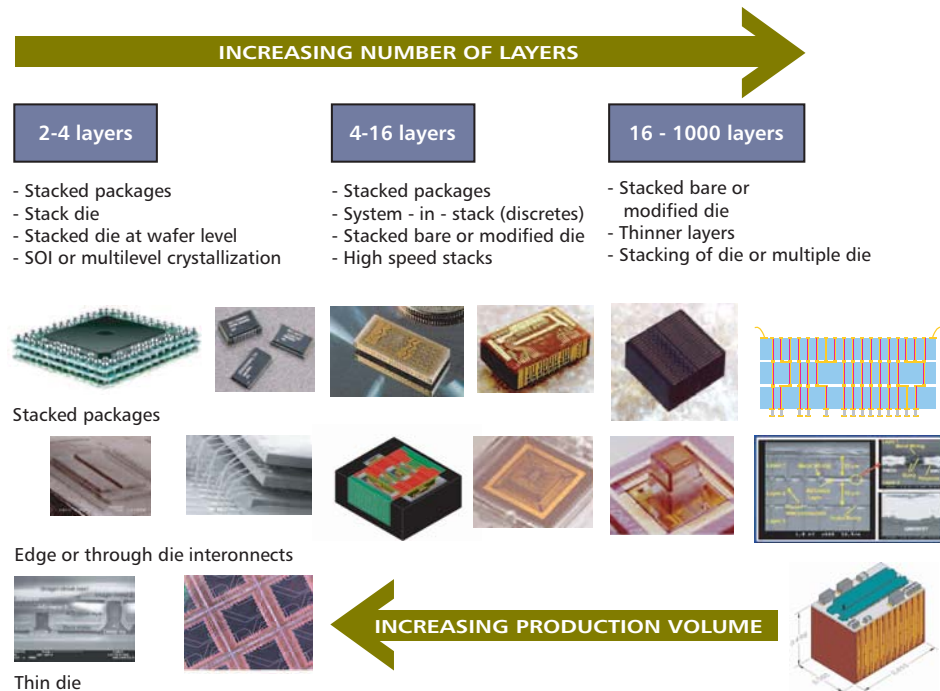


Figure 6-4. Categorization of 3D packaging approaches.

Superconducting electronic circuits based on NbN and Nb Josephson junctions (JJs) were inserted in stack layers and operated at temperatures as low as 4 K, indicating the flexibility and reliability of the material system used in 3-D chip stacks. Stacked systems were operated at 20 GHz. When the material selection and system design are judiciously performed, chip stacks have been shown to operate reliably with mean time to failure exceeding 10 years, in a wide temperature range from – 270 C to 165 C and in hostile environments subjected to 20,000 G. 3-D packaging provides an excellent alternative to satisfy the needs of the high functionality system and sub-system integration applications and can yield system architectures that cannot be realized otherwise.

When the power budget increases, thermal management layers can be inserted in the stack as alternating layers in addition to active layers. Experimental and analytical results indicated that up to 80W can be dissipated in a stack volume of 1cm³. Thermal resistance within the stack can be as low as 0.1 C/W.

6.2.2 3-D PACKAGING – READINESS

The design of 3-D packaged systems is technically feasible and fairly well understood. However, the design for larger applications and high-speed operation above 10 GHz have not been completely investigated.

6.2.3. 3-D PACKAGING – ISSUES AND CONCERNS

Minimal thermal resistance is critical for superconducting electronic applications where the components need to operate at 4 K. Other trade-offs associated with 3-D packaging are:

- Cost.
- Test.
- Reparability.

Testing of the complex 3D system-in-stack requires further development and better understanding of the system design. Added complexity will result from the high clock speed of SCE-based systems.

Another major issue (similar to secondary packaging) is the availability of a foundry. A dedicated and independent packaging foundry is critically needed for SCE-based applications. It is assumed that this foundry can be co-located with the MCM foundry.

6.2.4 3-D PACKAGING – ROADMAP AND FUNDING

A roadmap for the development of 3D packaging for SCE is shown in the figure below. The funding profile needed to maintain development pace is listed in Table 6-3. The funding assumes that the foundry is co-located with the MCM foundry. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)

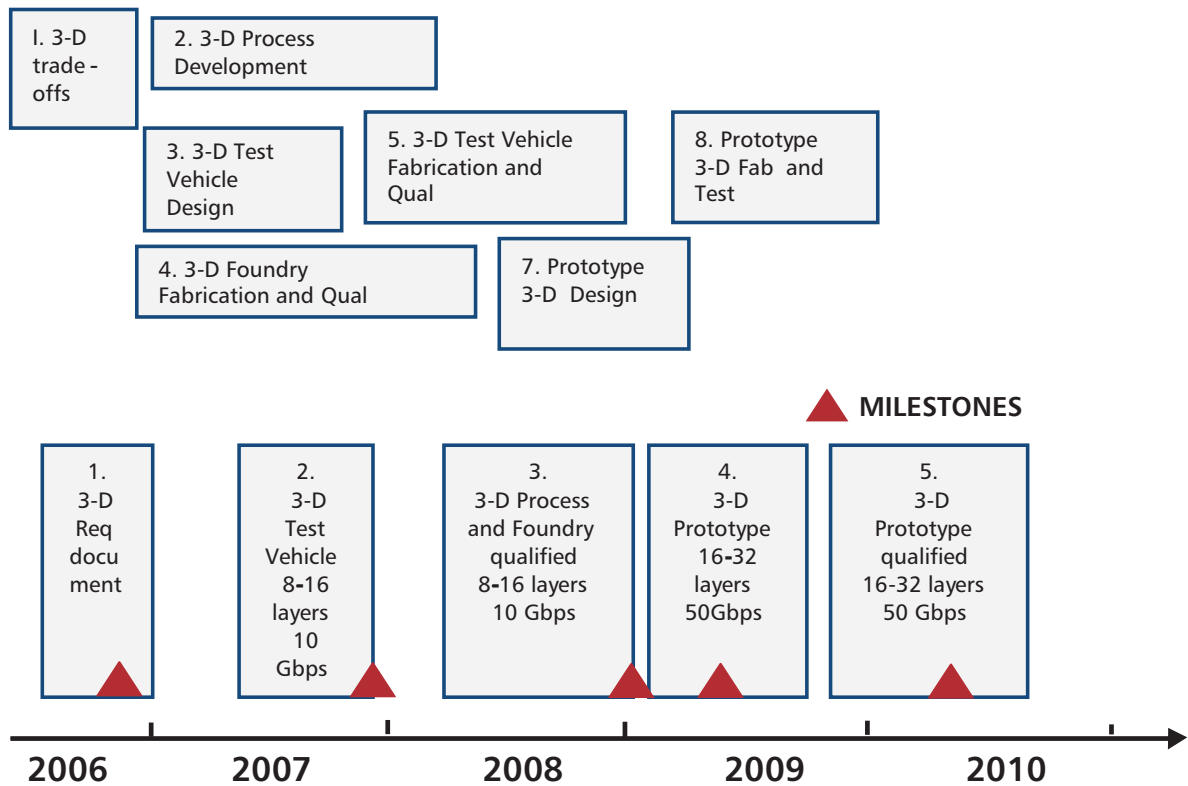


TABLE 6-3. 3D PACKAGING DEVELOPMENT COSTS (\$M)						
Year	2006	2007	2008	2009	2010	Total
Fabrication Equipment (\$M)	(covered in MCM Foundry investment)					
3D Development Total	1.3	2.0	2.0	0	0	5.3
Total Investment	1.3	2.0	2.0	0	0	5.3

6.3 ENCLOSURES AND SHIELDS

Enclosures and shields (illustrated in Figure 6-5) are required to operate SCE at cold temperatures, under vacuum or without magnetic interference, undue radiative heat transfer or vibration. The vacuum dewar is an integral part of the cryocooler operation and should be designed concurrently. A hard vacuum is needed to minimize convective heat transfer from the cold to warmer parts. Vibration of the system can translate into parasitic magnetic currents that may disrupt the operation of SCE circuits.




Figure 6-5. The construction details of a typical enclosure designed for operation at 4 K⁴. The overall enclosure is shown at left. The details of the housing for SCE circuits are shown in the cut-away at right.

6.3.1 ENCLOSURES AND SHIELDS – STATUS

Technical challenges for enclosures and shields include how to penetrate the vacuum walls and shields with large numbers of cables, both for high DC current power supplies and very-high-frequency (50-100 GHz) digital signals, while minimizing heat flow into the low temperature environment.

⁴ "Integration of Cryo-cooled Superconducting Analog-to-Digital Converter" D. Gupta et al, ASC 03'

Several companies have demonstrated the enclosures and shields for SCE circuits as prototypes or as products. Although other rack-mounted designs can be envisioned for distributed systems, most of them shared certain common characteristics. They:

- Were single-rack or single-stage standalone systems.
- Used a circularly symmetrical packaging approach.
- Were about 20-100 cm in size.

Packaging information about non-U.S. applications is scarce. Although several sources, such as NEC, report on their SCE circuit design and performance, they seldom report on their packaging approach. Furthermore, even if they did, the test results may only be for laboratory configurations rather than industrially robust ones.

6.3.2 ENCLOSURES AND SHIELDS – READINESS

The design of enclosures and shielding for SCE systems is technically feasible and fairly well understood. However, an experimental iteration of the design of larger applications (e.g., a petaflops-scale supercomputer where dimensions are in the order of several meters) has never been completed.

6.3.3 ENCLOSURES AND SHIELDS – PROJECTIONS

TABLE 6-4. PROJECTIONS FOR ENCLOSURES AND SHIELDING FOR SCE SYSTEMS				
Year	2004	2006	2008	2010
Type of enclosures and shields	Vacuum	Vacuum and Magnetic	Vacuum and Magnetic	Vacuum and Magnetic
Penetrating leads	500	1,000	5,000	10,000

6.3.4 ENCLOSURES AND SHIELDS – ISSUES AND CONCERNS

Due to the lack of a previous design at large scale (e.g., a functional teraflop SCE supercomputer), other potential issues may be hidden and can only be discovered when such an engineering exercise takes place. The past funding (commercial and government) for the development of such systems was in disconnected increments that pre-ented the accumulation of engineering know-how. Therefore, substantial learning and more development is needed to insure a functional SCE-based supercomputer system.

6.3.5 ENCLOSURES AND SHIELDS – ROADMAP AND FUNDING

A roadmap for the development of enclosures and shields for an SCE-based, large-scale computer is shown in the figure below. The funding profile needed to maintain development pace is listed in Table 6-5. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)

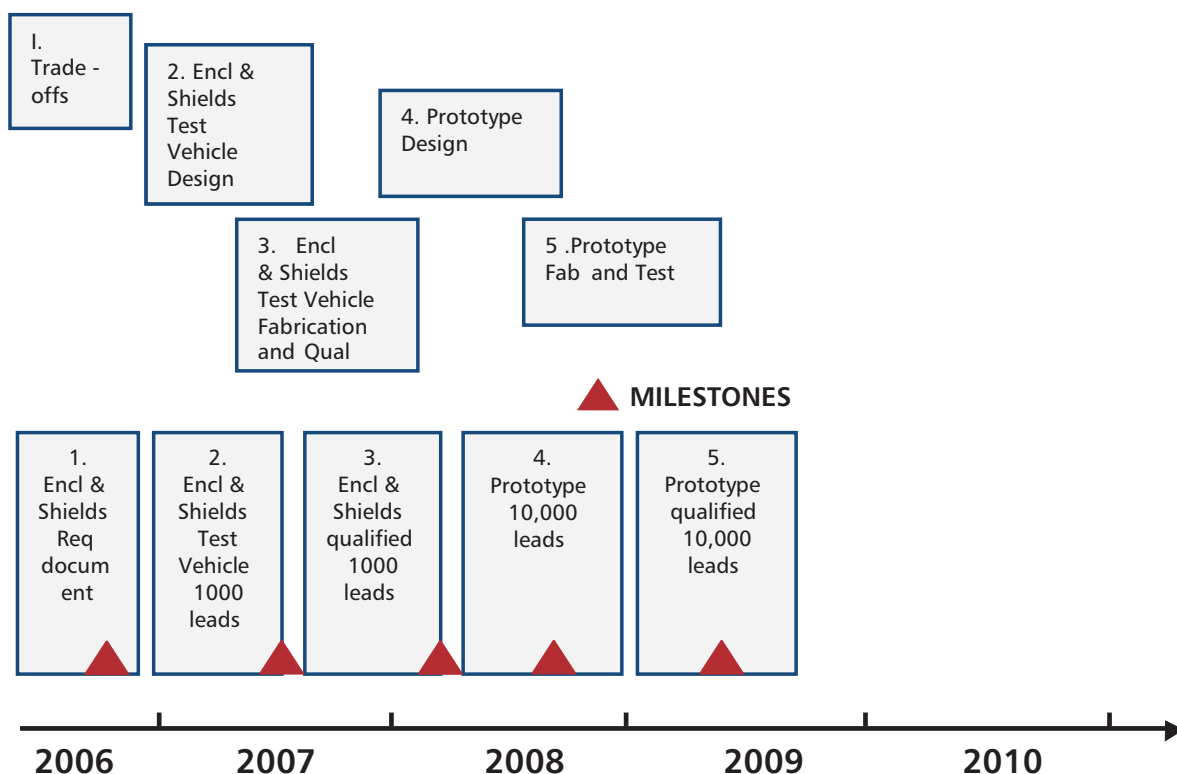


TABLE 6-5. ENCLOSURES AND SHIELDS DEVELOPMENT COSTS (\$M)

Year	2006	2007	2008	2009	2010	Total
Enclosures and Shields Development	2.0	1.8	1.3	1.3	1.0	7.4
Total Investment	2.0	1.8	1.3	1.3	1.0	7.4

6.4 COOLING

An SCE circuit-based system needs to operate at temperatures ranging from 4 to 77 K. Small, closed-cycle coolers commonly use helium gas as a working fluid to cool cold plates, to which the circuits are mounted in vacuum.

In selecting the cooling approach for systems, the heat load at the lowest temperatures is a critical factor. A large-scale, SCE-based computer with a single, central processing volume and cylindrically symmetric temperature gradient was analyzed under the HTMT program⁵. The heat load at 4 K arises from both the SCE circuits and the cabling to and from room temperature. The cable heat load (estimated to be kW) was the larger component, due to the large volume of data flow to the warm memory. Circuit heat load may be as small as several hundred watts. If all this heat at 4 K were extracted via LHe immersion, a heat load of 1 kW would require a 1400 liter/hour gas throughput rate.

It is possible to partition a large cryogenic system into many modules, each separately cooled. Approaches with a small number of large modules or approaches with many, smaller modules are feasible. Coolers for large-scale or small-scale modules are reasonable, with the larger ones being more efficient and actually more technologically mature than those of intermediate capacities. However, multiple smaller modules allow continued processing while some modules are taken off-line for repair and maintenance.

Since there are availability differences between large-scale and small-scale cryo-coolers, the status for each type is presented separately⁶.

6.4.1 COOLING – STATUS

Small Cryocoolers

Among commercial small cryocoolers, the GM coolers have been around the longest; tens of thousands of them have been sold. With the use of newer regenerator materials, two-stage GM coolers now can go down to temperatures as low as 2.5 K. GM coolers have a moving piston in the cold head and, as a result, produce vibration accelerations in 0.1 g range. They are quite reliable if scheduled maintenance of both the compressor units and the cold heads is performed regularly. A typical maintenance interval is about 5,000 hours for the compressor and about 10,000 hours for the cold head.

In the hybrid cooler and 77 K communities, there are a number of companies selling Stirling cycle coolers (e.g., Sunpower and Thales), but there are no commercial 4 K purely Stirling coolers on the market. These units tend to be physically smaller than GM machines of similar lift capacity due to the absence of an oil separation system. The vibration characteristics of Stirling machines are about the same as for GM machines. At this time—and under conditions where maintenance is feasible—commercial Stirling machines do not appear to be more reliable than GM machines.

Pulse tube coolers are commercially available from a number of vendors including Cryomech, Sumitomo, and Thales. Two-stage pulse tube coolers can produce over 1 W of cooling at 4.2 K on the second stage of a two-stage cooler and simultaneously produce about 50-80 W at 50 K. A pulse tube machine will inherently produce lower vibration accelerations (<0.003 g) on the cold head than a GM cooler. Since there are no moving parts in the cold head, pulse tube machines are easier to make very reliable. Maintenance intervals of 25,000 hours are recommended.

There are many companies involved in creating space cryocooler technology including Lockheed-Martin, Ball Aerospace, Creare, ETA Incorporated, Honeywell, Hughes, Mainstream Engineering Corporation, Mitchell/Stirling Machine Systems, Northrop Grumman (including TRW), Raytheon, Ricor, Swales Incorporated. The development was funded by DoD components, and the results are not commercially available.

⁵ HTMT Program Phase III Final Report, 2002

⁶ "Integration of Cryogenic Cooling with Superconducting Computers", M. Green, Report for NSA panel.

Large Cryocoolers

There are two European manufacturers of large machines, Linde in Germany and Air Liquid in France. There are a few Japanese manufacturers, but their track record is not as good as the Europeans. Dresden University in Germany is one of the few places in the world doing research on efficient refrigeration cycles. A system capable of handling kW level heat loads cools portions of the Tevatron accelerator at Fermilab in Illinois. The conceptual design of a single large system has been initiated under the HTMT program (Figure 6.6).




Figure 6-6. Concept for a large-scale system including cryogenic cooling unit for supercomputers⁷.

6.4.2 COOLING – READINESS

The technology for the refrigeration plant needed to cool large systems, along with its associated mechanical and civil infrastructure, is understood well enough to allow us to make technical and cost estimates. Small space and commercial cryocoolers are available, but engineering changes are needed for use in large-scale systems. Larger units and their components have been made in much smaller numbers and may require further development.

6.4.3 COOLING – ISSUES AND CONCERNS

Vibration can be a serious problem for superconducting electronics if the total magnetic field is not uniform and reproducible. The cryocooler itself can be mounted with flex mounts on the cryostat vacuum vessel, which has more mass and is, therefore, more stable than the cooled device.

The reliability issues can be mitigated by using smaller cryocooler units around a central large-scale cryogenic cooling unit. This approach adds modularity to the system and allows for local repair and maintenance while keeping the system running. However, the design issues resulting from the management of many cryocoolers working together are not yet well explored.

⁷ HTMT Program Phase III Final Report, 2002

Another issue is the cost of the refrigeration. Buildings and infrastructure may be a major cost factor for large amounts of refrigeration. Effort should be made to reduce the refrigeration at 4 K and at higher temperatures since the cost (infrastructure and ownership) is direct function of the amount of heat to be removed.

One key issue is the availability of manufacturers. The largest manufacturer of 4 K coolers is Sumitomo in Japan, which has bought out all of its competitors except for Cryomech. No American company has made a large helium refrigerator or liquefier in the last 10 years; the industrial capacity to manufacture large helium plants ended with the Superconducting-Super-Collider in 1993. Development funding may be needed for U.S. companies to insure that reliable domestic coolers will be available in the future. Development of an intermediate-sized cooler (between small laboratory-scale devices and large helium liquefiers) would be desirable to enable a system with multiple processor modules.

6.4.4 COOLING – ROADMAP AND FUNDING

A roadmap for the development of cryocoolers for an SCE-based, large-scale computer is shown in figure below. The funding profile needed to maintain development pace is listed in the accompanying table. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)




TABLE 6-6. COOLER DEVELOPMENT COSTS (\$M)						
Year	2006	2007	2008	2009	2010	Total
Cooler Development	1.3	0.9	1.1	0.6	0.6	4.5
Total Investment	1.3	0.9	1.1	0.6	0.6	4.5

6.5 POWER DISTRIBUTION AND CABLES

Superconducting processor chips are expected to dissipate very little power. The cryocooler heat load for a petaflops system will be dominated by heat conduction through input/output (I/O) and power lines running between low and room-temperature environments. To reduce heat load, it would be desirable to make the lines as small as possible in a cross section. However, requirements for large DC power supply currents and low-loss, high bandwidth I/O signaling both translate into a need for a large metal cross section in the cabling for low signal losses and low Joule heating. Therefore, each I/O design must be customized to find the right balance between thermal and electrical properties.

SCE circuits for supercomputing applications are based on RSFQ circuits that are DC powered. Due to the low voltage (mV level), the total current to be supplied is in the range of few Amperes for small-scale systems and can be easily kilo-Amperes for large-scale systems. Serial distribution of DC current to small blocks of logic has been demonstrated, and this will need to be accomplished on a larger scale in order to produce a system with thousands of chips. However, the overhead of current-supply reduction techniques on-chip can be expected to drive the demand for current supply into the cryostat as high as can be reasonably supported by cabling.

In addition, because petaflops systems will have very high I/O rates, Radio Frequency (RF) cabling, which can support high line counts serving thousands of processors with high signal integrity, is needed. Reliable cable-attach techniques for thousands of connections also require cost efficient assembly procedures with high yield.

6.5.1 POWER DISTRIBUTION AND CABLES – STATUS

Supplying DC current to all of the SCE chips in parallel would result in a total current of many kiloAmps. Several methods may be used to reduce this total current and heat load. One technique is to supply DC current to the SCE chips in series, rather than in parallel (known as current recycling). This technique of providing DC current to RSFQ has been demonstrated, but there is real estate and junction count overhead associated with this method. This overhead will drive system design to find a high DC current “comfort zone” for the power supply.

Another solution is to use switching power supplies. High voltages/low currents can be brought near SCE circuits and conversion to low voltages/high currents, all at DC, can occur at the point of use. However, this method employs high power field-effect transistor switches, which themselves can dissipate significant power.

If sufficient cooling power is available at the intermediate cryogenic temperature (such as 77 K), then high temperature superconductor (HTS) cables may be used to bring in DC current from 77 K to the 4 K environment. High temperature stranded DC cabling is a well known product, but there has been little demonstration of flexible ribbon cabling in HTS—the most desirable implementation for a large system such as a petaflops computer—where the component count and number of connections makes reliability, modularity, and assembly very difficult with individual cable.

For serial I/O in systems with low I/O counts, coaxial cables can be used for both high-speed and medium-speed lines. These cables are made in sections, with a middle section having a short length of stainless steel having a high thermal resistance, and with a bottom section of non-magnetic Copper for penetration into the chip housing.

For systems with hundreds or thousands of I/O lines, coaxial cables are not practical. To meet this challenge, flexible ribbon cables have been developed⁸ (Figure 6-7). These cables consist of two or three layers of copper metallization separated by dielectric films, typically polyimide. With three copper layers, the outer two layers serve as ground planes and the inner layer forms the signal lines, creating a stripline configuration. Stripline cables provide excellent shielding for the signal lines. Successful high-reliability and low cost cable-attach procedures for signals up to 3 GHz have been demonstrated, and this technique should allow operation up to 20GHz after minor modifications. Present flex cable line density is sufficient for the MCM-to-MCM connections, but is an order of magnitude lower than required for the external I/O.

⁸ “Cryogenic Packaging for Multi-GHz Electronics” T. Tighe et al, IEEE Tran. Applied Superconductivity, Vol 9 (2), pp3173-3176, 1999.

Some of the thermal and electrical issues associated with RF cabling can be mitigated by using fiber optics to carry information into the cryostat, which reduces parts count and thermal load. Using wavelength division multiplexing (WDM), each fiber can replace 100 wires. Optical fibers have lower optimal thermal conductivity than metal wires of the same signal capacity. A factor of 1,000 reduction in direct thermal load is achieved, because each fiber has 10 times less thermal conductance relative to wire cable. Although using optical means to bring data into the cryostat has been sufficiently demonstrated to be low risk, these advantages apply only to data traveling into the cryostat. There are known low-power techniques to convert photons to electrical current, but the reverse is not true. The small signal strength of SFQ complicates optical signal outputs, and modulation techniques at multi-Gbps data rates are unproven.

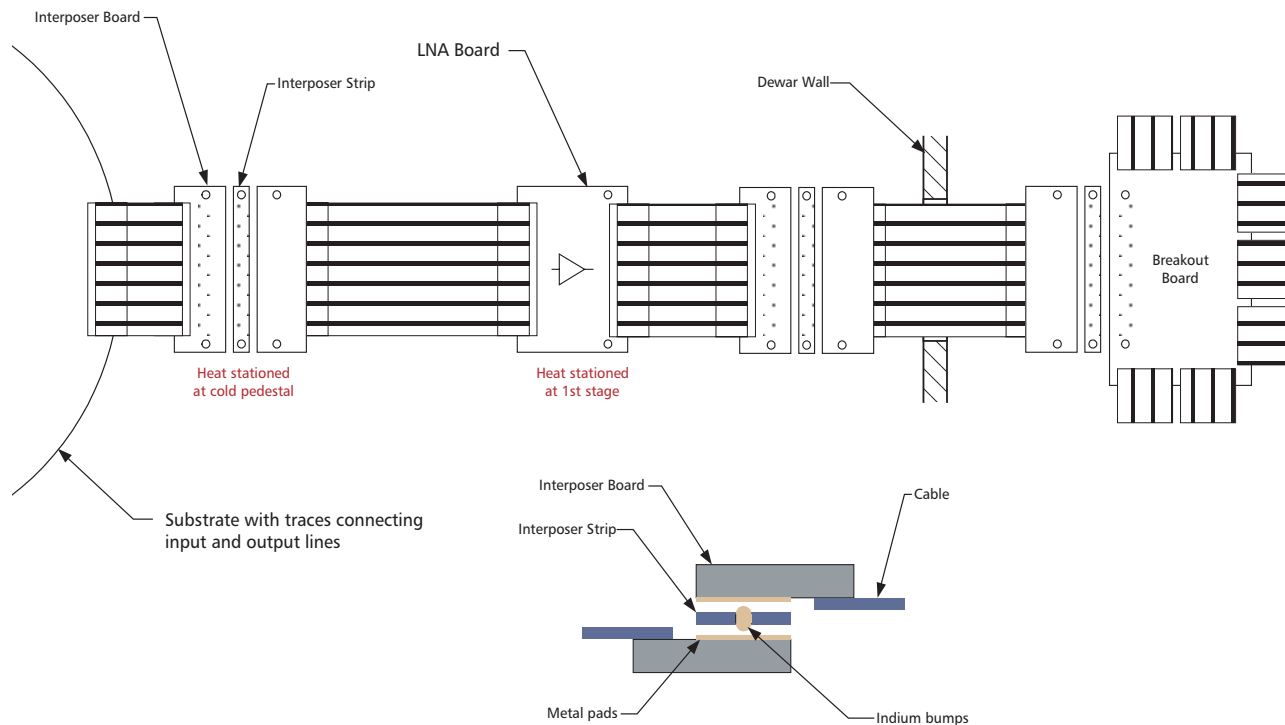


Figure 6-7. A high-speed flexible ribbon cable designed for modular attachment (Ref: NGST).

An alternate material to metals for the RF electrical cabling is high-temperature superconductors (HTS). HTS cables could be used between 4 K and 77 K. Conductors made from HTS materials, in theory, transmit electrical signals with little attenuation and with little heat load. This combination, which cannot be matched by non-superconductor metals, makes their use attractive. However, the high T_c materials are inherently brittle in bulk form and so lack the flexibility necessary to the assembly of systems with the 1,000's of leads expected in a supercomputer. The second generation magnet wire products based on YBCO thin films on a strong and flexible backing tape now coming into commercial production are being modified in an Office of Naval Research award for use as flexible DC leads. Carrying 10's to 100's of Amps from 77 K down to 4 K with low thermal loading should be quite straightforward. Given the large material difficulties of the high T_c materials, a better choice for the RF leads would be MgB_2 , which was recognized to be a 39K, s-wave superconductor only late in 2001. A scalable manufacturing process has been demonstrated at Superconductor Technologies that could today manufacture ten 1 cm x 5 cm tapes using the current substrate heater and could be straight-forwardly scaled up to a process that simultaneously deposits over a 10-inch diameter area or onto a long continuous tape.

6.5.2 POWER DISTRIBUTION AND CABLES – READINESS

Solution of the power distribution and cabling issues is technically feasible and fairly well understood at small-scale applications. However, the design of high-current DC cabling for larger applications, (e.g., a petaflops supercomputer where dimensions are in the order of several meters) remains to be done. For small systems, coaxial cabling has been the primary choice for RF cabling. Designs and data for cabling at ~10 GB/s and up to ~1,000 signal lines are available but must be developed further to support the many thousand RF lines required by a petaflops system. Extension to support line rates of 40 GHz or higher has not been fully explored. The mode of operation may change from microstrip/stripline to waveguides, which raises assembly and heat load issues. The good news is that, except for HTS cabling and RF cabling beyond 20 Gb/s, the commercial cable manufacturing industry should be able to support the needs of an RSFQ petaflops system with special order and custom products.

6.5.3 POWER DISTRIBUTION AND CABLES – PROJECTIONS

A projection of the current and projected state of suitable cabling for a petaflops RSFQ-based system is given below. Until a full architecture is developed, it is not possible to quantify the electrical, thermal, line count, and other requirements for the cabling.

TABLE 6-7. Current and Projected State of Suitable Cabling for Petaflops-scale System			
Year	2004	2007	2009
DC Power Cables	—	—	—
Lines	Single Conductor	Flex Cable	Flex Cable
Max DC current/line	~ 1 Amp	~100 mA	~500 mA
RF Signal Cables	—	—	—
Line Count	~100	~1,000	~2,000
Data Rate Supported	10 Gb/s	20 Gb/s	20 Gb/s
Current Supply	Direct, Parallel	Direct, Serial	Switched or Direct, Serial

6.5.4 POWER DISTRIBUTION AND CABLES – ISSUES AND CONCERNS

The operation at 50 GHz and beyond with acceptable bit error rates (BER) needs further development and testing, including physical issues (such as conductive and dielectric material selection, dielectric losses) and electrical issues (such as signaling types, drivers and receivers).

The combination of good electrical and poor thermal conductance is inherently difficult to meet since a good electrical conductor is a good thermal conductor. The required density of I/O lines in large-scale systems is a factor of 10 greater than the density achievable today (0.17 lines/mil) for flexible ribbon cables.

Optical interconnects may require the use of elements such as prisms or gratings. These elements can be relatively large and may not fit within the planned size of the cryostat. The power required for the optical receivers—or the thermal energy delivered by the photons themselves—may offset any gains from using the low thermal conductance glass fibers. A detailed design is needed to establish these trade-offs.

The use of HTS wires would include two significant technical risks:

- Currently, the best HTS films are deposited epitaxially onto substrates such as lanthanum-aluminate or magnesium oxide. These substrates are brittle, and have relatively large thermal conductance which offset the “zero” thermal conductivity of the HTS conductors.
- Any HTS cable would need to have at least two superconductor layers (a ground plane and a signal line layer) in order to carry multi-Gbps data.

Presently, multi-layer HTS circuits can only be made on a relatively small scale due to pin-holes and other such defects that short circuit the two layers together.

6.5.5 POWER DISTRIBUTION AND CABLES – ROADMAP AND FUNDING

A roadmap for the development of cables for an SCE-based, large-scale computer is shown below. The funding profile needed to maintain development pace is listed in Table 6-8. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)

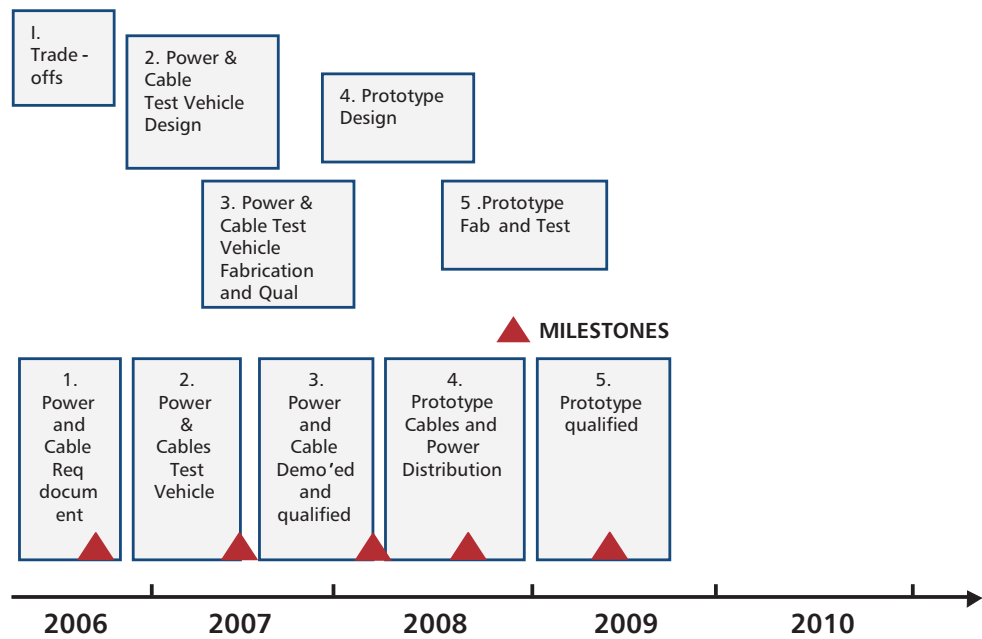


TABLE 6-8. CABLES AND POWER DISTRIBUTION DEVELOPMENT COSTS (\$M)						
Year	2006	2007	2008	2009	2010	Total
Cables and Power Distribution Development	2.0	2.3	3.3	3.5	3.6	14.7
Total Investment	2.0	2.3	3.3	3.5	3.6	14.7

6.6. SYSTEM INTEGRITY AND TESTING

Because a petaflops-scale superconducting supercomputer is a very complex system, offering major challenges from a system integrity viewpoint, a hierarchical and modular testing approach is needed. The use of hybrid technologies—including superconducting components, optical components and conventional electronic components, and system interfaces with different physical, electrical and mechanical properties—further complicates the system testing. This area requires substantial development and funding to insure a fully functional petaflops-scale system.

6.6.1 SYSTEM INTEGRITY AND TESTING – STATUS

System-level testing for conventional supercomputers is proprietary to a few companies such as Cray, IBM, Sun, and NEC. Testing of superconducting supercomputers has not yet been addressed. A limited amount of laboratory-level testing of components and sub-modules has been conducted at companies such as NGST, HYPRES, and NEC. Testing mostly addressed modular approaches at cold temperatures by providing standard physical interfaces and limited functional testing at high frequencies.

6.6.2 SYSTEM INTEGRITY AND TESTING – READINESS

The readiness for system-level testing is not well understood and requires further development and funding.

6.6.3 SYSTEM INTEGRITY AND TESTING – ISSUES AND CONCERNS

As stated above, testing of a large-scale superconducting system poses major challenges from engineering viewpoint. Key issues are:

- **Cold temperature and high temperature testing, probes and probe stations:** A manufacturable and automated high-speed testing of SCE circuits operating at 50 GHz clock frequencies requires considerable development in terms of test fixtures and approaches. Issues such as test coverage versus parametric testing, use of high frequency, and high frequency at cold temperatures for production quantities needs to be further developed. The engineering efforts needed are considerable.
- **Assembly testing:** Testing becomes further complicated when SCE circuits are grouped at MCM and/or in 3-D modules. Test engineering know-how at this level is minimal. Approaches such as built-in self test can be utilized but require considerable development effort including fixturing and support equipment issues along with simple but critical questions such as “what to test” to insure module level reliability. Testing becomes more challenging when these types of assemblies include other interfaces such as optical I/Os.
- **Production testing:** Repeated testing of chips and assemblies require automated test equipment. Such equipment does not exist even for conventional high-speed electronic systems.

6.6.4 SYSTEM INTEGRITY AND TESTING – ROADMAP AND FUNDING

A roadmap for the development of test approaches and testing equipment for an SCE-based, large-scale computer is shown in the figure below. The funding profile needed to maintain development pace is listed in Table 6-9. More details are presented in Appendix L: Multi-Chip Modules and Boards. (The full text of this appendix can be found on the CD accompanying this report.)

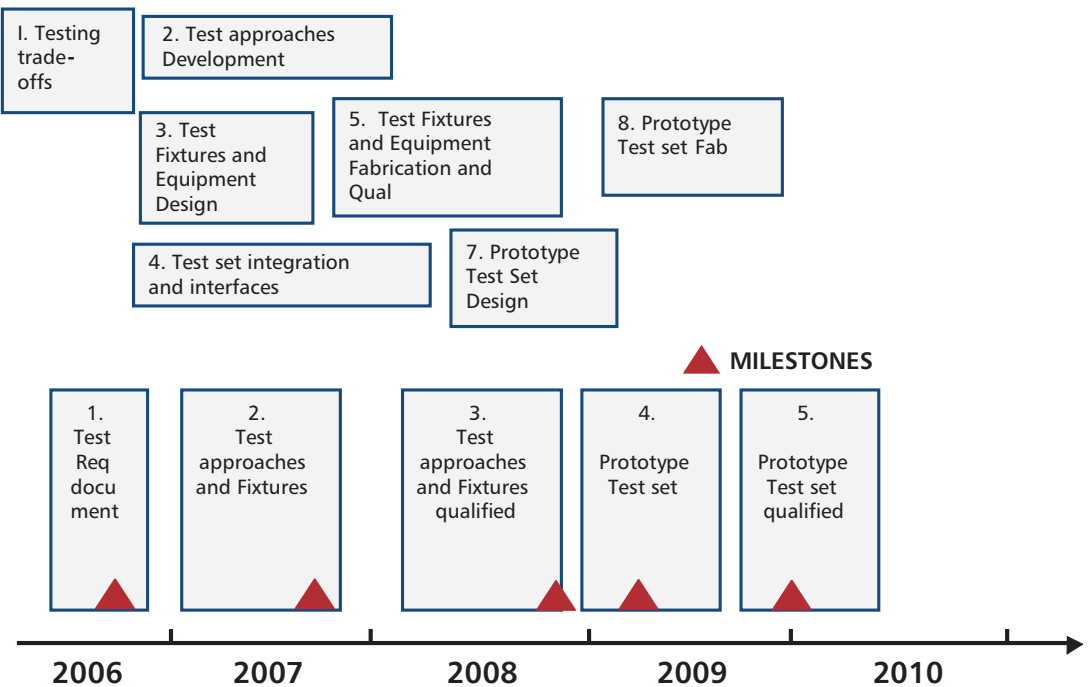


TABLE 6-9. SYSTEM TESTING COSTS (\$M)

Year	2006	2007	2008	2009	2010	Total
Equipment Total	1.0	2.0	2.0	0.0	0.0	5.0
Test Development Total	1.3	1.8	1.8	1.9	1.9	8.7
Total Investment	2.3	3.8	3.8	1.9	1.9	13.7

Appendix A

TERMS OF REFERENCE

Date: 10 September 2004

Terms of Reference: OCA Superconducting Technology Assessment

Task

With the approval of the Director, NSA, the Office of Corporate Assessments (OCA) will conduct an assessment of superconducting technology as a significant follow-on to silicon for component use in high-performance computing (HEC) systems available after 2010. The assessment will:

- Examine current projections of:
 - material science.
 - device technology.
 - circuit design.
 - manufacturability.
 - general commercial availability of superconducting technology over the balance of the decade.
- Identify programs in place or needed to advance commercialization of superconducting technology if warranted by technology projections.
- Identify strategic partnerships essential to the foregoing.

First-order estimates of the cost and complexity of government intervention in technology evolution will be needed. The assessment will not directly investigate potential HEC architectures or related non-superconducting technologies required by high-end computers other than those elements essential to the superconducting technology projections.

Background

Complementary metal oxide semiconductor (CMOS) devices underpin all modern HEC systems. Moore's Law (doubling of transistor count every 18 months) has resulted in steady increases in microprocessor performance. However, silicon has a finite life as devices shrink in feature size. At 90 nanometers, there are increasing difficulties experienced in material science, and clock and power features in commodity microprocessors, and already signs are appearing that the major commodity device industry is turning in other directions. Most companies are planning on fielding lower power devices with multiple cores on a die (with increased performance coming from device parallelism vice clock/power increases). It appears doubtful that commodity microprocessors will shrink much beyond the 65 nanometer point. While the Silicon Industry Association (SIA) roadmap projects silicon usage well into the next decade, it will almost certainly come from increased parallelism, not speed. Looking ahead, the SIA projects superconducting Rapid Single Flux Quantum (RSFQ) technologies as a promising replacement for systems requiring substantial increases in processor speeds.

Specific Tasks

- Assess the current state of superconducting technologies as they apply to microprocessors, hybrid memory-processor devices, routers, crossbars, memories and other components used in general purpose HEC architectures and other high speed telecommunications or processing applications.
- Where possible, identify and validate the basis for SIA projections in current superconducting technology roadmaps. Seek expert opinion where necessary.

- Construct a detailed superconducting technology roadmap to reach operational deployment, to include:
 - a. All critical components and their technical characteristics at each milestone.
 - b. Dependencies on other technologies, where such exist.
 - c. Significant developments/actions that must occur (e.g., creation of fabrication facilities).
 - d. Continuing research needed in parallel with development of technology, in support of the roadmap.
 - e. Critical, related issues tied to particular milestones.
 - f. Estimated costs, facilities, manpower etc. tied to particular milestones.
- The roadmap will be developed at three levels:
 - a. An aggressive, full government funding level with availability of industrial strength technology by 2010.
 - b. A moderate government funding level with availability of industrial strength technology by 2014.
 - c. Non-government funding (industry reliance) level.
- The study will conclude by presenting options for government decision makers and the expected outcome of choosing each option in respect to superconducting technologies.
- The study will be fully documented and completed within 6 months of initiation, but not later than 31 March 2005.

Assessment Structure

The study will be conducted primarily by a team of outside experts augmented by Agency experts in the field. These outside experts will be drawn from industry, academia and other government agencies. Team membership is attached. OCA will provide funding, logistic and administrative support to the study. The study, along with recommendations for follow-on actions, will be forwarded by OCA to the Director and Deputy Director for review and approval.

Appendix B

PANEL MEMBERS

INDUSTRY/ACADEMIA	
John T. Pinkston (Chairman)	Dept. of Computer Science and Electrical Engineering University of Maryland, Baltimore County
Ted Van Duzer	Dept. of Electrical Engineering and Computing Sciences University of California, Berkeley
Arnold Silver	Subject Matter Expert (SME)
John X. Przybysz	Northrop-Grumman
Sheila Vaidya	Lawrence Livermore National Laboratory
Hirsch Mandelberg	Science Applications International Corporation
John Carson	Irvine Sensors
Volkan Ozguz	Irvine Sensors
John Spargo	Northrop-Grumman
Deborah Van Vechten	Program Officer for Superconducting Electronics Office of Naval Research, Department of Navy
Loring Craymer	Jet Propulsion Laboratory
Alan Kleinsasser	Jet Propulsion Laboratory
Mikhail Dorojevets	Dept. of Electrical and Computer Engineering Stony Brook University, New York

GOVERNMENT	
George Cotter	Director for Information Technology NSA/CSS Chief Information Officer
Nancy Welker	STA Senior Corporate Advisor Chief Technical Officer Operational Test Authority
Doc Bedard	STA Senior Technical Advisor Technical Director
Mike Escavage	STA Technical Support & Coordination Chief, Technology Test & Evaluation
Jim Pinkham	STA Webmaster, System Administrator
James Rupp	STA Technical Support
Stanley Hall	STA Technical Support
Sylvia Tether	STA Administration, Office Manager

Appendix C

GLOSSARY

TERMS/DEFINITIONS	
BER	Bit Error Rate
CDR	Critical Design Review
CMOS	Complementary Metal Oxide Semiconductor
CNET	Cryostatic NETwork
CRAM	Cryogenic RAM
DARPA	Defense Advanced Research Projects Agency
DoD	Department of Defense
DRAM	Dynamic RAM. Has relatively slow access time—in the tens of nanoseconds—but high memory density.
FLOPS	Floating Point Operations Per Second. Used as a unit of performance.
FLUX	Not an acronym. This is the name of a series of superconducting electronics VLSI experiments.
Gb	Gigabit
Gbps	Gigabit per second
GB	GigaByte
GFLOPS	GigaFLOPS
GHz	Giga Hertz—one billion times per second
HEC	High-End Computer
HECRTF	High-End Computing Revitalization Task Force
HPCS	High Productivity Computing Systems
HRAM	Holographic RAM
HTMT	Hybrid Technology Multi-Threaded architecture
IBM	International Business Machines
IC	Integrated circuit or “chip”

TERMS/DEFINITIONS	
IHEC	Integrated High End Computing; a congressionally-mandated R&D Plan
LLNL	Lawrence Livermore National Laboratory
JJ	Josephson junction
JPL	Jet Propulsion Laboratory
JTL	Josephson Transmission Line
kW	kilowatt(s)
MCM	Multi-Chip Module
mW	milliwatts(s)
MW	megawatt(s)
MOS	Metal Oxide Semiconductor
MRAM	Magnetoresistive Random Access Memory
NASA	National Aeronautics and Space Administration
NGST	Northrop Grumman Space Technology
NSA	National Security Agency
NSF	National Science Foundation
NRZ	Non-Return to Zero
peta	Prefix meaning 10^{15}
PCB	Printed Circuit Board
PFLOPS	PetaFLOPS
PIM	Processor In Memory. A hardware architecture in which both processing logic and memory are placed on the same IC. The principle advantage is to allow direct access to a row (usually 2048 bits) at a time.
RAM	Random Access Memory
RF	Radio Frequency
RSFQ	Rapid Single Flux Quantum
R&D	Research and Development
SCE	Superconducting Electronics
SFQ	Single Flux Quantum

TERMS/DEFINITIONS	
SIA	Silicon Industry Association
SOA	Semiconductor Optical Amplifier.
SRAM	Static RAM. Can be quickly accessed in a small number of nanoseconds, but has relatively low memory density.
SUNY	State University of New York (at Stony Brook)
tera	Prefix meaning one trillion (10^{12})
TFLOPS	TeraFLOPS
TRW	TRW, Inc. (Now NGST)
VLSI	Very Large Scale Integrated (circuits)
W	Watt

Appendix D

INTRODUCTION TO SUPERCONDUCTOR SINGLE FLUX QUANTUM CIRCUITRY

The Josephson junction (JJ) is the intrinsic switching device in superconductor electronics. Structurally a JJ is a thin-film sandwich of two superconducting films separated by a very thin, non-superconducting, barrier material. For the junctions of interest here, the superconductors are Nb and the barrier material is a dielectric, aluminum oxide. The barrier is sufficiently thin, $\sim 1\text{nm}$, that both normal electrons and superconducting electron pairs can tunnel through the barrier.

Electrically, a tunnel junction between normal metals is equivalent to an ohmic resistor shunted by a capacitor. If the metals are superconductors, however, electron pairs can tunnel at the same rate as normal electrons. Only pairs can tunnel at zero voltage, while both pairs and normal electrons tunnel if a voltage appears across the junction. Pair tunneling gives rise to a third parallel channel, resulting in unique electrodynamics and highly-nonlinear current-voltage characteristics.




Figure 1. Equivalent circuit of Josephson junctions. I_c represents the nonlinear switch.

JJs are characterized by the equivalent circuit shown in Figure 1, where I_c is the critical current (maximum current at zero voltage) of the junction, C is the capacitance, and R_d is the resistance. C is capacitance of the junction.

R_d is the parallel combination of the internal junction resistance and any external shunt. Low-inductance shunts are commonly employed in RSFQ circuits to control junction properties. The intrinsic junction resistance is very voltage dependent as shown in Figure 2a. It is very large compared with typical external shunts for voltages below the energy gap, $V_g \sim 2.7\text{ mV}$ for niobium (Nb), and very small at the gap voltage. Above the gap voltage, the resistance becomes ohmic, with a value R_N , the resistance of the junction in the non-superconducting state. The product $I_c R_N$ is approximately V_g .

The time-dependent behavior of this equivalent circuit is given by summing the three current components illustrated in Figure 1, resulting in a non-linear differential equation,

$$(1) \quad I = I_c \sin \int V dt + \frac{V}{R} + C \frac{dV}{dt}.$$

The first term represents a parametric Josephson inductance whose value scales as $L_J = \Phi_0/2\pi I_c$. Damping is controlled by the parameter $\beta_C = R^2 C/L_J$, the ratio of the junction time constants RC and L_J/R .

In digital circuits, JJs operate in two different modes: voltage-state (latching) and single flux quantum (non-latching). Figure 2 illustrates the static current-voltage characteristics of these two device modes, which are also characterized as under-damped and over-damped, respectively. JJs are always current-driven at zero-voltage, so their behavior depends on their response to the external current. In Figure 2a, the I-V characteristics are multi-valued and hysteretic, such that the junction switches from $V = 0$ to V_0 at $I = I_c$. If the current is reduced near zero, the junction resets in the zero-voltage state. This provides a two-state logic voltage that was the basis of the IBM and Japanese computing projects in the 1970's and 1980's. If the junction is shunted by a small resistor, the I-V characteristic is single-valued as shown in Figure 2b. The voltage can increase continuously from zero as the current increases above I_c .




Figure 2a. DC electrical characteristics of voltage-state latching junctions




Figure 2b. DC electrical characteristics of SFQ non-latching junctions

The early work exemplified by the IBM and the Japanese Josephson computer projects exclusively used voltage-state logic, where the junction switching is hysteretic from the zero-voltage to the voltage state. This necessitated an AC power system in order to reset the junction into the zero-voltage state. The speed of this technology was limited to about 1 GHz.

Single flux quantum is the latest generation of superconductor devices and circuits. The I-V curve for SFQ operation is single-valued and the devices are DC powered. A fundamental property of JJs in the voltage state is that the junction produces precisely reproducible SFQ pulses at a frequency proportional to the voltage:

$$f_{\text{SFQ}} = V / \Phi_0 \quad (2)$$

where $\Phi_0 = 2.07 \text{ mV}\cdot\text{ps}$ is the magnetic flux quantum. Each pulse represents one quantum of magnetic flux, $2.07 \times 10^{-15} \text{ Webers}$, passing through the junction. At $100 \text{ }\mu\text{V}$, the SFQ frequency is 50 GHz . Thus, invisible in the DC characteristics, the junction DC voltage is the result of generating identical SFQ pulses according to Eq. (2). A 2 ps pulse is approximately 1 mV . In SFQ circuits, each switching junction is associated with a small inductor L that can compress and store a flux quantum. A parameter $\beta_L \sim 1$ defines the relation between β_L and I_C ,

$$\beta_L = \frac{2\pi L I_C}{\Phi_0} = \frac{L}{L_J}; \quad (3)$$

Switching time is a critical factor for digital applications; the minimum pulse width and maximum frequency are limited by parameters J_C , I_C , C' , and R , where C' is the specific capacitance of the junction and R is generally an external shunt resistance. SFQ junctions are designed for optimal speed (i.e., near critical damping). An external shunt resistor is used to insure that $\beta_C \sim 1$. Then, the SFQ pulse width is:

$$\tau = \frac{\Phi_0}{2\pi I_C R} = \sqrt{\frac{\Phi_0 C'}{2\pi J_C}}. \quad (4)$$

Thus, the maximum operating frequency scales as $J_C^{1/2}$. Figure 3 shows that the measured speed of asynchronous flip-flops, the simplest SFQ logic circuit, follows this rule.




Figure 3. Measured speed of static dividers varies as $J_C^{1/2}$.

RSFQ electronics is faster and dissipates less power than the earlier superconductor logic families that were based on the voltage-latching states of JJs. Furthermore, SFQ circuitry has capitalized on IC processing and CAD tools already developed for the semiconductor industry to achieve very much higher performance (faster at lower power) than semiconductors with the same generation fabrication tools.

In the early 1990's, Prof. K.K. Likharev and his team relocated from Moscow State University to Stony Brook University. Since that time, the U.S. has led the development of single flux quantum electronics. These efforts were driven in part by the Department of Defense University Research Initiative on Superconducting Digital Electronics, but also included significant contributions from HYPRES, NIST, Northrop Grumman, TRW and Westinghouse. Japan has recently embarked on a project to develop Nb SFQ digital technology for high data rate communications.

Nb IC technology employing Nb-AlO_x-Nb tunnel junctions has been adopted almost universally. This technology operates successfully below 4.5 K. The intrinsic ability to produce complex SFQ circuits at any given time is limited by the available infrastructure such as IC fabrication, design, packaging, and testing tools. The fact that IC fabrication for SFQ chips is generations behind semiconductor fabrication is the result of limited funding, not because of fundamental limitations of this technology. Nevertheless, SFQ development has produced significant small circuit demonstrations directed at high-speed signal processing applications such as A/D converters. A common metric for SFQ digital device capability in any given superconductor IC fabrication process is the speed of asynchronous toggle flip-flops (FF), which has reached as high as 770 GHz, as shown in Figure 3.

Some of the important features of SFQ circuits for high-end computing are:

- Fast, low power switching devices (JJs) that generate identical single flux quantum pulses.
- Lossless superconducting wiring for power distribution.
- Latches that store a magnetic flux quantum, $\Phi_0 = 2 \times 10^{-15}$ volts-sec.
- Low loss, low dispersion superconducting transmission lines that support "ballistic" data and clock transfer.
- Cryogenic operating temperatures that reduce thermal noise and enable low power operation.

When the junction energy, $\Phi_0 I_C / 2\pi$, is compared with kT at 4 K, we find that the equivalent thermal noise current, $I_n = 2\pi kT / \Phi_0 = 168$ nA. The critical currents in practical circuits must be sufficiently larger than I_n to achieve the desired BER. In today's technologies, minimum I_C is about 100 μ A, nearly 600 times larger than I_n .




Figure 4. Passive transmission lines can propagate picosecond SFQ pulses without dispersion at the speed of light in the line. Velocity in the stripline line is $\sim c/3$.

SFQ logic operates by generating, storing, and transmitting identical SFQ pulses. Data pulses are transmitted by direct connection between adjacent gates. Data is transmitted to distant gates as SFQ pulses through impedance-matched passive microstripline or stripline as shown in Figure 4. When an SFQ pulse is fed into the left-most inductor, a clock-wise current $i = \Phi_0/L$ appears in the inductor-JJ pair, which adds to the current in the junction. If the sum of the SFQ current and the bias exceeds I_C , the junction switches and the SFQ pulse is transmitted down the line.

Figure 5 illustrates three basic SFQ circuit configurations: data latch, OR gate, and AND gate. The two-junction comparator is the basic decision-making element in SFQ logic (Figure 5a). The data latch of Figure 5a includes an input junction (J1), an inductor to store an SFQ pulse ("1"), and a comparator J2/J3 that determines whether or not an SFQ pulse will be transmitted for each clock pulse. An SFQ pulse appearing at J1 will switch J1 and store one flux quantum in the inductor between J1 and J2/J3. The stored flux quantum adds a current $\sim \Phi_0/L$ in J3. If there is a flux quantum in the latch when the clock pulse arrives, J3 switches, an SFQ pulse is transmitted, and the latch is reset to 0. If there is no flux quantum in the latch when the clock arrives, the current in J3 is insufficient to switch J3 and no SFQ pulse is transmitted. The OR and AND gates depicted in Figure 5 represent the inputs to the latch. In the OR gate, an SFQ pulse at either input will be transmitted to the output and into a latch. As usual, SFQ pulses have to arrive "simultaneously" at both inputs to transmit an output in Figure 5c. The two inputs would be clocked from preceding latches to ensure simultaneous arrival.




Figure 5. Representative SFQ gates

(a) SFQ data latch (DFF)

(b) Or gate (Merger)

(c) AND gate

The comparator must be robust in order to withstand parameter variations and thermal noise that can speed-up or delay switching. These have the effect of reducing operating margins, which translates into higher incidence of errors. The effective operating margins are those for which an acceptable bit-error-rate (BER) is achieved. This is quantified in terms of the acceptable range of DC current injected between the two junctions. The minimum junction critical current (I_C) is a compromise between thermal noise considerations discussed above and existing lithography limits for junctions and inductors. Since junctions are connected to small inductances, L , and $L I_C \sim 2 \times 10^{-15}$ Wb, the maximum local inductance is ~ 20 pH. An increase in I_C must be accompanied by an equivalent reduction in junction area. At 20 kA/cm^2 , the smallest junctions will be $\sim 0.8 \text{ } \mu\text{m}$. To achieve a factor of two increase in switching speed, I_C should be 80 kA/cm^2 and the smallest junctions would be about $0.4 \text{ } \mu\text{m}$. Lower I_C requires larger inductors that occupy more space and reduce gate density. Higher I_C increases the effect of parasitics.

Bias resistors are used to convert low voltage supplies to fixed current sources for each junction. They are not shown in any of the circuits, but are implied in series with the supply current. Although power dissipation is very low compared with all other technologies, there is continuous power dissipation in the bias resistors (R_B) equal to $0.5 I_C^2 R_B$, assuming junctions are biased at $\sim 0.7 I_C$. A more convenient measure is $0.7 I_C V_B$, since R_B is set equal to $V_B / 0.7 I_C$. For 2 mV and $100 \text{ } \mu\text{A}$, the static power dissipation is $\sim 0.14 \text{ } \mu\text{W/JJ}$. Assuming 10 JJs/gate , the static power per gate is $\sim 1.4 \text{ } \mu\text{W}$.

Similar to CMOS, the irreducible power dissipation only occurs during SFQ switching events when a voltage appears across the junction. Assuming a junction switches at the frequency f , the minimum power dissipation is $P_{\text{SFQ}} = I_C V = I_C \Phi_0 f$. Based on the present minimum $I_C \sim 100 \mu\text{A}$, $P_{\text{SFQ}} = 0.2 \text{ nW/GHz}$ per junction. Assuming ten JJs per gate and five JJs switch each cycle per gate, $P_{\text{SFQ}} = 1 \text{ nW/GHz}$ per gate. At 100 GHz, $P_{\text{SFQ}} = 0.1 \mu\text{W/gate}$. The ratio of the static power to the switching power is ~ 20 at 100 GHz.

Additional information about JJ technology can be found in numerous books, including:

- T. Van Duzer and C.W. Turner, *Principles of Superconductive Devices and Circuits*, (Elsevier, NY, 1981).
- Alan Kadin, *Introduction to Superconducting Circuits* (John Wiley and Sons, NY, 1999).

Review articles on RSFQ technology include:

- K.K. Likharev and V.K. Semenov, "RSFQ Logic/Memory Family: A New Josephson Junction Technology for Sub-Terahertz Clock-Frequency Digital Systems," *IEEE Transactions on Applied Superconductivity*, vol. 1, pp. 2-28, March 1991.
- K.K. Likharev, "Superconductor Devices for Ultrafast Computing," in *Applications of Superconductivity*, H. Weinstock, Ed., Dordrecht, Kluwer, 1999.
- A.H. Silver, A.W. Kleinsasser, G.L. Kerber Q.P. Herr, M. Dorojevets, P. Bunyk, and L. Abelson, "Development of superconductor electronics technology for high-end computing," *Superconductor Science and Technology*, vol. 16, pp. 1368-1374, December 2003.

Appendix E

SOME APPLICATIONS FOR RSFQ

The primary drawback of RSFQ for widespread acceptance is the requirement for cryostatic operation. Despite this, there are a number of applications for which high speed processing is essential, including communications applications which already use low-temperature sensors and receivers.

The ever increasing need for network security makes it desirable to have processing at or near the network switch level (i. e., delays are unacceptable) so that content scanning becomes feasible without compromising network performance. In a commercial context, very high serial processor speed is desirable to allow cellular phone CDMA-based networks to sort out additional message streams from the composite signal in the digital domain, thereby increasing the system capacity and reducing costs.

For the DoD in general, the increasing tempo of warfighting and demands for minimal collateral damage with no waste of expensive munitions puts increasing pressure on communications systems, as do requirements for intelligence gathering and processing in a world full of mobile communications devices and other electronics. Low thermal noise is often desirable for sensors, including radio receivers; digital superconducting electronics can provide processing capability at 4 K so that raw data can be processed into actionable knowledge locally without wasting precious communications bandwidth and transmitter power on useless data. Software-defined radios, wideband dynamic bandwidth allocation, and receivers that use correlation to provide exactly matched filtering become feasible with sufficient digital processing capability.

Data from NRL suggests that superconducting electronics devices are radiation tolerant, with an upset rate on the order of 1/10,000th that of hardened CMOS. Coupled with the low-power operation characteristic of RSFQ, digital superconducting electronics appears to be an attractive technology for spaceborne applications, as well.

Supercomputing Applications and Considerations

The Integrated High End Computing (IHEC) report documented several areas in which there is an ever-increasing need for high performance computing:

- Comprehensive aerospace vehicle design.
- Signals intelligence (processing and analysis).
- Operational weather/ocean forecasting.
- Stealthy ship design.
- Nuclear weapons stockpile stewardship.
- Multi-spectral signal and image processing.
- Army future combat systems.
- Electromagnetic weapons development.
- Geospatial intelligence.
- Threat weapon systems characterization.

The report further identified four bottlenecks (in no particular order) suffered by these applications:

- Memory performance (latency/bandwidth/size).
- CPU performance.
- Programming productivity.
- I/O system performance (internal communications; also storage and external I/O).

Computational performance depends not just on CPU performance, but also memory and inter-node communications. For large-scale supercomputing, there is an additional constraint: the amount of parallelism that can be extracted from an application. Modern supercomputers are composed of a large number of processors operating concurrently. Each processor supports one or more threads, where a thread is a part of a software program that may execute independently of the rest of the program; multiple threads may execute concurrently. Whenever the amount of processing parallelism in an application falls below the number of threads supported by the hardware, throughput falls below peak for the system on which the application is being run. Some applications have high degrees of extractable parallelism, but many do not. Relative to CMOS, RSFQ needs more than an order of magnitude fewer hardware threads to provide the same level of computational throughput.

Within a parallel computer, the performance of communications between nodes can be viewed as a mix of performance for near-neighbor and random inter-node communications and does not scale linearly with node count. The link bandwidth consumed by a message depends on the number of “hops” required to reach its destination: a one-hop message consumes half the link bandwidth that a two-hop message does. Bandwidth consumed by near-neighbor communications roughly scales as the number of nodes while random inter-node message bandwidth consumption does not. For Clos networks and hypercubes, required system bandwidth for randomly addressed messages scales as $N \log N$ (N the number of nodes); for tori, the scale factor is kN^2 , where k is a constant that depends on the number of dimensions and the length of each dimension. Minimizing node count also minimizes the cost of the inter-node communications subsystem.

It is difficult to estimate the practical limits for RSFQ-based supercomputers. RSFQ has a potential 100x advantage over CMOS; this might translate to a limit of a few hundred PFLOPS—about 1,000x what has been achieved to date with CMOS. At this point, however, the limiting factor may not be processor performance, but rather the memory and storage required for a balanced machine. The rule-of-thumb for memory is that a balanced machine should have 1 Byte/FLOPS; disk storage should be on the order of 10-100x memory, with still greater amounts of tape storage.

DoD Communications Applications

In addition to the use of superconductive digital logic for extreme computing power, in the 2010-2015 time frame RSFQ has the potential to deliver the RF system hardware that enables the increased operational tempo, fully collaborative operations, and extremely minor collateral damage assumed in the Network Centric Warfare vision.

Three classes of systems that enable the delivery of these goals are:

- Software defined radios to provide interoperability.
- Low distortion/high power efficiency transmitters to allow single units to handle multiple simultaneous signals.
- Simultaneous scan-and-stare receivers utilizing matched filtering to provide operational awareness.

Each class will utilize RSFQ logic's extreme clock speed to allow manipulation of the true time dependence of the signal in the digital domain. In addition, a significant source of error and noise in the operation of mixed signal components is removed by the exact quantum mechanical reproducibility of each RSFQ pulse. The >10x lower thermal noise in the circuits (due to their 4 K operating temperature and lower characteristic impedance) also facilitates utilizing the extreme sensitivity to magnetic fields in setting the minimum detectable signal. The high speed system clock allows time averaging between base band changes in the signal information content.

The simultaneous scan-and-stare receivers could significantly improve our battlespace (and homeland) situational awareness by both accessing more simultaneous signals by parallelizing only the digital filters and by processing the signals in closer to real time. Especially in single antenna systems, this is enabled by being able to digitize accurately wide swaths of frequency and then digitally selecting the specific sub-band and bandwidth desired for each signal without losing accuracy. This compares favorably to selecting the signals of interest in the analog domain (via heterodyning techniques) prior to digitization. Only RSFQ logic has demonstrated this ability. Cross-correlation techniques on the RF waveform turn out to be a straightforward way of implementing matched filtering under real time software control and harvesting additional processing gain in comparison to base band approaches.

Software radios are exemplified by the Joint Tactical Radio System (JTRS) program. The goal is to unify the hardware required to receive and transmit all legacy waveforms and facilitate the introduction of new, higher-data-rate waveforms. The idea is that the software running at a given time will determine how the hardware functions. Inter-banding (the essential enabler of interoperability) is achieved by using different waveform software on receive and transmit. So far, the conventional approaches implemented in semiconductor technologies have not been highly successful. For example, JTRS cluster 1 faces a major re-evaluation and potential termination after EOA in the spring of 2005. They are having trouble breaking away from simply co-locating multiple radios in a federated design. The above discussed uniquely demonstrated ability of RSFQ logic to implement true digital reception allows one to change waveforms by changing the control parameters in the digital filters that operate directly on the RF signals. In receivers, this direct reception eliminates many expensive analog components and their associated spurious signals and drastically simplifies the processing.

In transmitters, the RSFQ clock speed allows linearization to be done straightforwardly using the carrier and its harmonics, not some deep subharmonics that cannot capture the subtleties of the actual signal. Indeed, many systems, including JTRS, want one transmitter to do multiple tasks, ideally simultaneously. However, this is very difficult to do with reasonable energy efficiency with today's current hardware power amplifiers – they are typically highly non-linear when operated in their highest energy efficiency mode (often below 40% even for a single tone). To avoid transmitting large numbers of substantial amplitude spurious signals due to the non-linearity, separate amplifiers for each signal and power combiners that throw away 50% of energy at each combine are used. By enabling predistortion in the digital domain to compensate for the amplifier non-linearity, these combiners can be eliminated. Better signal quality and substantially better energy efficiency – especially important whenever a limited fuel supply or lifetime of battery are concerns – are expected to result from the use of RSFQ logic in such systems.

Commercial Potential

Three applications with commercial potential were identified at the 2001 Workshop on Superconducting Electronics:

- A digital signal processor for use in CDMA base stations for cellular telephone networks.
- A scanning firewall to catch malicious content.
- Low-power information servers.

Analog high temperature superconductor components operating at 55-80 K are already in use in commercial CDMA networks. They offer greatly enhanced rejection of out-of-band signals and a much reduced system (thermal) noise floor. Interference between users generally limits the capacity of any given base station and stringent signal strength control measures are in place to reduce the tendency for one caller to drown out the others. Digital signal processing, especially successive interference cancellation, has been implemented to filter out as much of this interference as possible. However, the processing must be done in real-time. The maximum speed of a CMOS digital filter of this type does not offer the increment over that of a CMOS digitizer with sufficient resolution to capture the total signal needed to deliver as many resolved users as commercially desirable. Digital filters in the tens of gigahertz such as are achievable in RSFQ logic should provide a solution. Indeed, there is an active program in Sweden to demonstrate this early application of digital superconducting electronics and its significant potential of return on investment. Since the wireless market has already accepted the use of cryo-cooled electronics in base stations, the shift to low-temperature superconductor technology will be less of an issue. Small commercial cryo-coolers are readily available at 4 K.

In the information assurance context, viruses, worms, and hacker break-ins have become all too common. Sensitive information needs to be protected from “prying eyes.” At the same time, the information needs to be readily available to authorized users, some local and some at remote locations. Like the wireless communications problem, CMOS processing would be unacceptably slow, but a superconducting solution is possible. The potential market for scanning firewalls is potentially very large: if available, these would be considered essential by any moderate to large organization.

The low-power information server is a primary goal of the current Japanese digital superconductor effort. The application core is a superconducting router, with additional superconducting logic to provide the necessary intelligence. Such a server could support upwards of 100,000 transactions per second. The Japanese find the low-power argument compelling; although the power argument is less compelling in the U.S., the information server market is quite large and expanding. Again, there seems to be sufficient market for this to be a commercially viable solution.

Spaceborne Applications

Near-Earth Surveillance and Communication

Near-earth applications come in at least three flavors:

- Providing the military with interoperability, message prioritization, and routing in otherwise optical high throughput communications networks among satellites.
- On-orbit data reduction.
- The ability to sense much weaker signals because of a reduction in thermal noise.

Unlike the commercial terrestrial communications and internet, there is no hardwired, high throughput backbone to military communications. Nor is it feasible to consider laying long-lived cables among mobile nodes. What we have is several disjointed systems of high altitude communications satellites. There are relatively advanced concept studies looking into using optical communications among these satellites to form a multi-node network of high capacity assets, the equivalent of the internet backbone. While an optical approach to the actual message traffic, the task of reading the message header and rerouting the signal to the next node—especially when it is necessary to resolve contention and implement message prioritization—is not currently within the optical domain’s abilities. Moreover, the electronics domain is required to interconvert one satellite system’s message formats into another’s in order to achieve interoperability and thereby to reduce the total system cost. The clock rates proposed for superconducting digital systems match those of optical communications and make SCE an ideal candidate for this application.

In imaging and surveillance systems, most of the data collected tells us that nothing has changed and may be discarded. Yet many systems are designed to ship all data to the ground for processing in a centralized location, often introducing a substantial lag time before the useful information is ready. This scheme wastes substantial bandwidth and transmission power, and reduces the value of the information. Instituting distributed, on-board processing—especially if software defined and having a “send-everything” option—can deliver the actionable knowledge contained in the data with better efficiency in terms of time, bandwidth, and power. Superconductive digital electronics, once matured, should be up to this task.

Deep Space Applications

The cold background temperatures of deep space make the use of superconductive electronics for the entire receive chain highly attractive. Superconductive antennas can set exceptionally low system noise temperatures. Superconductive mixers are already the work horse of astrophysical receivers above 100 GHz. And the low noise temperatures and extreme sensitivity possible in superconductive ADC and following digital filters allow weak signals to be sensed.

One long-time dream for NASA is a mission to Pluto. Because a lander mission does not appear feasible, a flyby of perhaps 30 minutes duration is the most likely scenario. With communication lag times of many hours between Earth and Pluto, the space probe would require fully autonomous data gathering and analysis capabilities. Analysis during flyby is critical for optimizing the quality of data collected. There is very little hope of providing the power needed for CMOS-based data processing, but superconducting electronics could provide the processing capabilities needed at a fraction of the power budget for CMOS.

NASA is actively pursuing a program for missions to the icy moons of Jupiter. Here, the problem for CMOS is both cold and intense radiation. However, CMOS can be shielded and radiation-hardened, and radioisotope thermoelectric generators provide both heat and electrical power. Nuclear propulsion is being developed as a key technology and will provide more electrical power. Given these workarounds, superconducting electronics is not critical for achieving mission goals; however, RSFQ could provide a boost in computing power and a higher-performance communications system. This might serve to increase the science return of such missions.

For many years, NASA has been carrying out missions to Mars. Recently, it has committed to an ambitious program of returning to the moon as a way-station for (potentially manned) travel to Mars. With all this activity, NASA has come to realize the need for an interplanetary communications network to provide the bandwidth needed for returning scientific data to Earth. The Mars Telecom Orbiter (MTO), expected to launch in 2009, is intended to provide some of that bandwidth. Just as superconducting electronics is attractive for use in cellular phone networks, so, too, it would be attractive in missions similar to the MTO. Additionally, the superconducting analog of MTO could provide high-performance computing capability for in situ analysis of data.

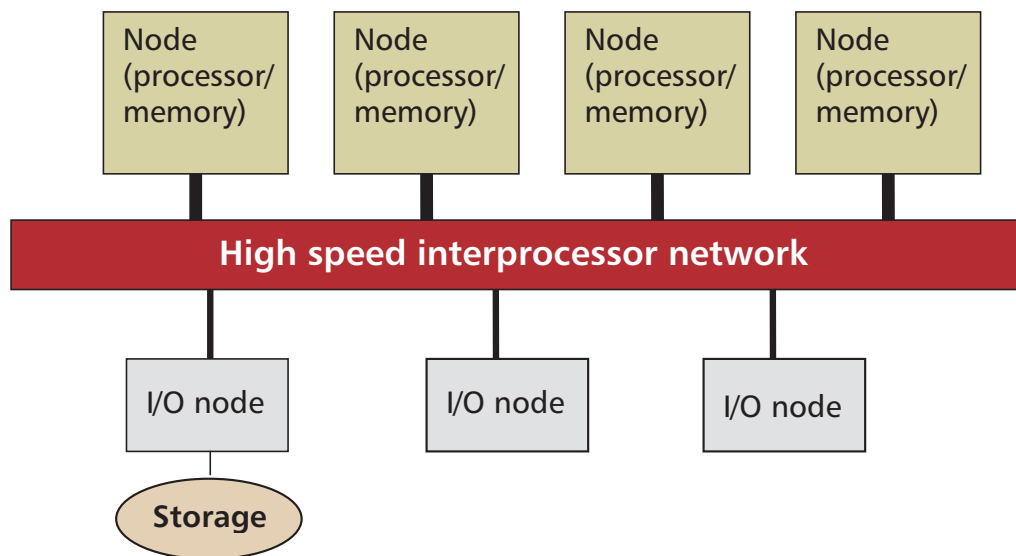
Appendix F

SYSTEM ARCHITECTURES

System-level Hardware Architecture

The primary challenges to achieving extremes in high performance computing are 1) integrating sufficient hardware to provide the necessary peak capabilities in operation performance, memory and storage capacity, and communications and I/O bandwidth, and 2) devising an architecture with support mechanisms to deliver efficient operation across a wide range of applications. For peta-scale systems, size, complexity, and power consumption become dominant constraints to delivering peak capabilities. Such systems also demand means of overcoming the sources of performance degradation and efficiency reduction including latency, overhead, contention, and starvation. Today, some of the largest parallel systems routinely experience floating point efficiency of below 10% for many applications and single digit efficiencies have been observed, even after efforts towards optimization.

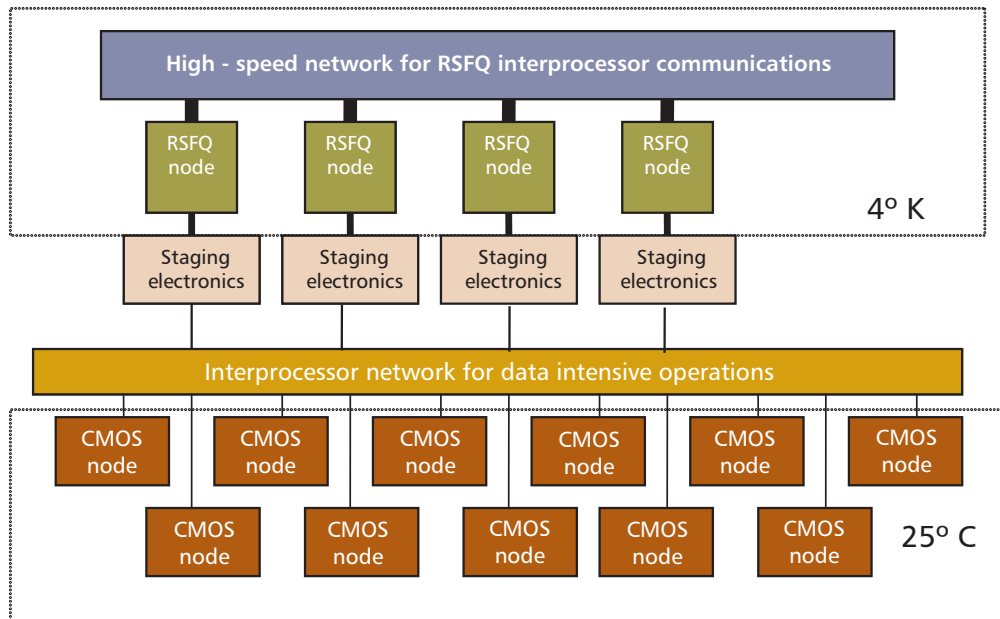
The “obvious” approach for deploying RSFQ in a high-end computer is to replicate a traditional MIMD (multiple instruction, multiple data) architecture. Processing nodes (processors and memory) are interconnected with a high-speed network:



Unfortunately, physical realities intervene and such a system would be unbalanced: it would have very high speed processing and high latency inter-processor communications. In addition, the differential between storage access latencies and processor speed would very high, on the order of 10^6 - 10^9 .

It would also require large amounts of memory per node (the rule of thumb usually applied is 1 byte/FLOPS); this in turn increases the physical size and thus the inter-node communications delays. This architecture supports compute-intensive applications, but would perform poorly for data-intensive applications. Indeed, considerable time would be wasted in loading memory from disk. Most supercomputing applications have a mix of compute-intensive and data-intensive operations; a different approach to architecture is needed.

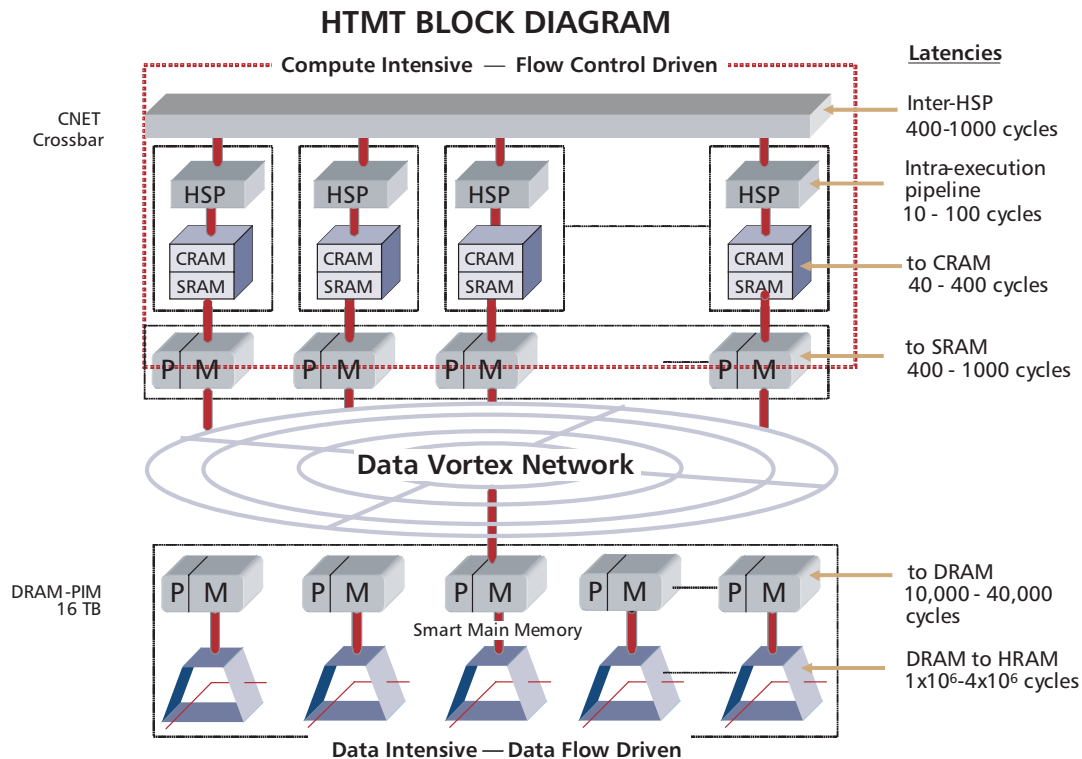
Only one significant effort to architect a balanced supercomputer architecture that could effectively incorporate RSFQ has been carried out to date. In the initial HTMT study, Sterling, Gao, Likharev, and Messina recognized that RSFQ provided the blinding processing speed desired for compute-intensive operation, while CMOS was needed to provide the high-density, low-cost components needed for data-intensive operation and access to conventional I/O and storage devices. This led to a three-tier processing architecture with two communications layers:



Not shown are the attached storage and I/O devices. These would be connected to CMOS nodes or to the data-intensive network and would achieve high aggregate bandwidth as a result of parallel access. A closer look at memory technologies led to inclusion of a layer of “bulk” memory—high-density, block-access memory such as might be provided by holographic storage or disk. Even with the high feature densities achieved by CMOS, a petabyte of memory takes up significant space.

(Note that partitioning the architecture into distinct compute-intensive and data-intensive segments correlates with the two operating temperature regimes. We found this useful in defining technology boundaries for this study.)

The HTMT project identified and pursued several technologies that might efficiently implement a design based on the above conceptual architecture. At its conclusion, the HTMT architecture had taken the form concurrently. In practice, there are limits to the degree of concurrency supportable by a given application at any one time. Low latency is desirable, but latency is limited by speed-of-light considerations, so the larger the system, the higher the latency between randomly selected nodes. As a result, applications must be capable of high degrees of parallelism to take advantage of physically large systems.



Since the demise of the HTMT project, processor-in-memory technology has advanced as a result of the HPCS program; the other technologies—RSFQ processors and memory, optical Data Vortex, and holographic storage—have languished due to limited funding.

Latency, Bandwidth, Parallelism

Attaining high performance in large parallel systems is a challenge. Communications must be balanced with computation; with insufficient bandwidth, nodes are often stalled while waiting for input. Excess bandwidth would be wasted, but this is rarely a practical problem: as the number of “hops” between nodes increases, so does the bandwidth consumed by each message. The result is that aggregate system bandwidth should increase not linearly with numbers of nodes but as $N \log N$ (Clos, hypercube) or N^2 (toroidal mesh) to sustain the same level of random node-to-node messages per node. Large systems tend to suffer from insufficient bandwidth from the typical application perspective.

Physical size is another limitation. Burton Smith likes to cite Little’s Law:

$$\text{latency} \times \text{bandwidth} = \text{concurrency}$$

in communications systems which transport messages from input to output without either creating or destroying them. High bandwidth contributes to high throughput: thus, high latencies are tolerated only if large numbers of messages can be generated and processed concurrently. In practice, there are limits to the degree of concurrency supportable by a given application at any one time. Low latency is desirable, but latency is limited by speed-of-light considerations, so the larger the system, the higher the latency between randomly selected nodes. As a result, applications must be capable of high degrees of parallelism to take advantage of physically large systems.

Execution Models

Organizing hardware for an RSFQ-based supercomputer is one thing, but how do we organize processing and software to maximize computational throughput? The widely used approach of manually structuring software to use message passing communications is far from optimal; many applications use synchronizing “barrier” calls that force processors to sit idle until all have entered the barrier. Similarly, the overlapping of computation with communications is often inefficient when done manually.

Execution models are the logical basis for maximizing throughput, while system hardware architecture provides the physical resources to support the execution model. Hardware architectures not tied to an execution model are unlikely to support optimal throughput; conversely, execution models that ignore hardware constraints are unlikely to be efficiently implemented. Mainstream parallel computers tend towards a message-passing, SPMD (single program, multiple data) execution model and hardware architecture with a homogeneous set of processing nodes connected by a regularly structured communications fabric. (The physical structure of the fabric is not visible to the programmer.) This execution model was developed to handle systems with tens to hundreds of processing nodes, but computational efficiencies decrease as systems grow to thousands of nodes. Much of the inefficiency appears due to communications fabrics with insufficient bandwidth.

There have been a number of machines developed to explore execution model concepts. Some of these—most notably the Thinking Machines CM-1 and the Tera MTA — were developed for commercial use; unfortunately, both of these machines were successful in achieving their technical goals but not their commercial ones.

The HTMT study is noteworthy in that the conceptual execution model was developed to match the hardware constraints described above, with the added notion that multi-threading was a critical element of an execution model. The conventional SPMD approach could not be mapped to the hardware constraints, nor could multi-threading alone define the execution model. The solution was to reverse the conventional memory access paradigm: instead of issuing a memory request and waiting for a response during a calculation, the HTMT execution model “pushes” all data and code needed for a calculation into the very limited memory of the high-speed processors in the form of “parcels.” Analytical studies of several applications have demonstrated that the execution model should give good computational throughput on a projected petaflops hardware configuration.

While little data is yet publicly available, both the IBM and Cray HPCS efforts appear to be co-developing execution models and hardware architectures. The IBM PERCS project is “aggressively pursuing hardware/software co-design” in order to contain ultimate system cost. The Cray Cascade project incorporates an innovative hardware architecture in which “lightweight” processors feed data to “heavyweight” processors; the execution model being developed has to effectively divide workload between the lightweight and heavyweight processors. Extensive use of simulation and modeling is being carried out to minimize the risk associated with adopting an innovative architecture.

While there is some research into execution models for distributed execution (Mobile Agent models, for example), there is little research on execution models for large-scale supercomputers. The Japanese supercomputing efforts (NEC and Fujitsu) have focused on engineering excellence rather than architecture innovation. The GRAPE sequence of special-purpose supercomputers uses an execution model designed to optimize a single equation on which N-body simulations depend.

Software Considerations

Program execution models are implemented in both hardware and software. To effectively evaluate a novel system architecture, it is necessary to map applications onto the execution model and develop an understanding of how the software runtime system would be implemented. If applications cannot be expected to exhibit high performance on the execution model, there is little point in attempting to build a supercomputer that would implement the model.

Software has been—and continues to be—a stumbling block for high-end computing. Developers have had to cope with SPMD (a program is replicated across a set of nodes, each of which processes its own data interspersed with data transfers between nodes) execution models that force them to explicitly deal with inter-processor communications and the distribution of data to processing nodes. Distributed memory systems have required message passing between nodes; shared memory systems have required explicit synchronization to control access to shared memory. Either way, it may take as much time to develop and test the communications infrastructure as it does to develop and test the core processing logic in an application.

In order to lessen the pain of communications programming in parallel applications, the communications APIs hide knowledge of the underlying hardware topology from applications. While this does contribute to ease of programming, it also effectively prevents the application developer from constructing an optimal mapping of the application to the underlying hardware. This is unfortunate, given the fact that bandwidth is limited, and that both bandwidth utilization and message latencies are strongly affected by the mapping of the application to the hardware.

New approaches to software tools and programming models are needed to take advantage of the execution models developed for RSFQ-based supercomputers. Further, these tools are best developed in parallel with the hardware and system architectures to provide feedback during that development. This approach reduces the risks associated with systems that are both innovative and costly.

The Tera (now Cray, Inc.) MTA provided an excellent example of this approach. By directly supporting multiple threads per processor in hardware, threads became a ubiquitous feature of the programming model, and synchronization (the MTA is a logically shared-memory architecture) was necessarily supported at the hardware level. Since these features were ubiquitous, they were supported by the compilers and other elements of the development tool chain. The compiler support was good enough that developers could provide “hints” to the compiler for tuning parallel applications instead of having to develop (possibly inefficient) communications infrastructures for their applications.

Development Issues and Approaches

There are no technical barriers to development of an architecture for an RSFQ-based supercomputer. However, the novel hardware constraints will require an innovative hardware architecture and execution model; an ongoing program of architectural innovation, analysis, modeling, and simulation would mitigate the risks that will occur if a large supercomputer is built without first carefully developing an operational proof of concept.

The presence of multiple processor types with differing instruction sets in a system architecture will present a challenge for software development tools. Cray is addressing this for the Cascade project, but Cascade will not have the extremes of processor performance of an RSFQ-based supercomputer. Although more work is needed in this area, significant investment should be deferred until RSFQ technology development reaches the point at which the specific problems to be overcome in software become clear.

What can be done to ensure architectural readiness at the time RSFQ technology is sufficiently mature? Open-ended architectural research is probably not the answer, but HTMT demonstrated a synergy between architectural development and advances in device technologies—the architecture provided a clear goal for the device technologies, and the device technologies served to identify a clear set of constraints for the architect to satisfy. Some promising approaches are:

- Holding periodic study competitions to develop innovative system concepts. These might occur every two years, with the studies taking six months to a year. There are two examples of this approach: the 1996 NSF/NASA/DARPA competition and the HPCS Phase I studies. The 1996 competition had one clear winner—the HTMT architecture—but the runner-up proposal identified processor-in-memory as an important technology that was incorporated into the subsequent HTMT Phase II Project. The 1996 competition was between academic researchers. The HPCS Phase I competition demonstrated that industry also could foster innovation; however, the industry entries pushed less aggressive technology and architecture goals. At the same time, the industry competition identified several approaches as viable. An open competition might encourage both innovation and feasibility.
- Identifying architectural shortcomings and investigate alternatives as early as possible. Moderately funded architectural modeling and simulation efforts for candidate architectures could be pursued, with close involvement between the technology developers and the architectural modelers to identify implementation tradeoffs. This would include elaboration of execution models and the runtime software designs that might support such models.
- Developing programming environment and simulated runtime, with the major investment starting two years before expected technology maturation. Software is key to enabling system operation.
- Encouraging ongoing system engineering activities to identify gaps in the technology and projected componentry, as well as to provide coordinated oversight of the technology development.
- Application benchmarking. Early contributions would be analyses to support the architectural modeling efforts; demonstration in a simulation environment would begin as the technologies near maturation.

Appendix G

ISSUES AFFECTING RSFQ CIRCUITS

Margins and Bit-Error-Rates

Any successful digital circuit technology requires large margins and low bit-error-rates (BER). Margins are the acceptable variation for which the circuit performs as designed. BER is the error-incidence per device per operation. Achieving high margins at high speed in RSFQ circuit technology depends critically on the symbiotic marriage of circuit design and chip fabrication. Margins and, consequently, BER in superconductors are affected by many factors, including all parameter spreads and targeting, noise, gate design, power (current) distribution cross-talk, signal distribution cross-talk, ground plane return currents, moats and flux trapping, and timing jitter.

Bit Error Rate and Integration Scale

It is necessary to balance performance and manufacturing issues. IC values fall in the range 0.1-0.2 mA, to ensure adequate noise immunity while keeping the total current and power dissipation as small as possible. A high J_C value improves speed, but requires smaller junction size, so parameter spreads may be more difficult to control. Characteristic parameter spreads must be many times smaller than the operating margins of the circuits in order to achieve VLSI densities. The requirements for petaflops computing are especially restrictive in that a very low BER is required. If one assumes one error per year is acceptable, then a 64 bit processor with 100 gates/bit, capable of petaflops throughput, would require a BER $\sim 10^{-26}$. This BER can be used to quantify the requirements for foundry IC process control.




Figure 1. Calculated BER of a 2-JJ comparator as a function of operating margin, for $I_C=0.2$ mA (lower curve), and $I_C=0.1$ mA (upper curve). Horizontal line is for 10^{-26} BER, which corresponds to one error/year for petaflops capacity

The two-junction comparator, the basic logic element in most SFQ logic, must be sufficiently robust to withstand parameter variations as well as thermal noise. Parameter variations reduce operating margins, which result in higher thermally-induced error rates. The effective, or error-free, operating margins are those for which an acceptable BER is achieved. This can be quantified in terms of the margins on a bias current injected between the two junctions. Gates are typically designed with bias margins of approximately $\pm 0.3I_c$. Variations have the effect of shifting the comparator threshold from its designed value. Variations in the ratio of junction critical currents, I_{c1}/I_{c2} , and the current scale of the input SFQ pulse, given by product LI_c , are particularly relevant. The BER for this device (Figure 1) is well-approximated by:

$$BER = \frac{1}{2} \operatorname{erfc} \left[\sqrt{\frac{\Phi_0}{\sqrt{2} I_c 4 k_B T}} (I_x \pm 0.30 I_c) \right]$$

BER is independent of J_c , but dependent on temperature, T , and I_c . For BER of 10^{-26} , operating margins are $\pm 10\%$ for $I_c = 0.2$ mA. This is referred to as the noise-free margin. For $I_c = 0.1$ mA, the noise-free margin is vanishingly small at this BER.

The BER requirement quoted above appears to be feasible, at least in terms of fundamental limits. In practice, lower operating temperature, larger I_c , or both might be needed. However, use of fault-tolerant design techniques has the potential to dramatically relax BER requirements. Fault-tolerant techniques presently employed in multi-processor systems should be directly applicable to RSFQ petaflops-scale computing. This might include parity bits, especially critical in the integer pipeline and memory, and check-pointing. Other types of errors should be considered as well, such as hard failures. For example, if inoperative processors could be taken off-line individually, they would need to be repaired or replaced only during scheduled maintenance intervals. Such techniques could greatly improve machine reliability.

The size of the error-free margin, M , can be used to determine process control, σ , needed to achieve VLSI. If the yield of a single gate is P , then the yield of a circuit of N gates is given by:

$$p^N = \operatorname{erf}^N \left(M / \sqrt{2\sigma} \right) \approx 1 - N \operatorname{erfc} \left(M / \sqrt{2\sigma} \right)$$

To achieve appreciable yield at 100K gates/cm² with $M=10\%$, σ must be less than ~ 2 to 3%, as shown in Figure 2. This also indicates that integration scale is practically unlimited when process control reaches a certain threshold. Thus, parametric yield may not be the limiting factor. Other aspects of yield are discussed on the following page.




Figure 2. Yield as a function of process variations (σ) normalized to the circuit operating margin (M), for circuits ranging from 1 to 109 gates.

Gate speed

Junction SFQ switching speed has been demonstrated to increase as $J_C^{1/2}$ up to at least 50 kA/cm². For $J_C = 20$ kA/cm² in Nb, the SFQ pulse width is ~ 2 ps and the measured speed of a static divider is 450 GHz, in good agreement with prediction. Junction speed does not appear to be a limiting factor for 50 to 100 GHz clocking, where the clock period is 20 and 10 ps, respectively. J_C can be increased to 100 kA/cm² to achieve the limiting speed of the junction. For J_C well above 20 kA/cm², junction properties change and new physical simulation models are likely to be required. However, the physics of such junctions is already understood.

Gate margins

Margins are the acceptable variation in external parameters (supply bias) and internal parameters (I_C , L , R , ...) for which the circuit performs as designed. We need large bias margins that are centered at the same external bias for all gates. Shrinking margins have been reported for relatively large circuits ($>10^3$ JJ's), particularly as the clock rate has increased above 10 GHz. However, test results for such circuits have only recently been reported, so there is not adequate data to validate specific solutions.

Chip complexity

Greater chip complexity is achieved by:

- Reducing gate footprints. Factors that help shrink the gate footprint are gate design, reducing the footprint of inductors and resistors (vertical resistors and inductors would have the smallest footprint) of the same electrical value, directly grounded junctions, and strapped contacts. Junctions, particularly directly grounded, are a small part of the physical size of gates. Until we reach dimensions near 10 nm, inductance values scale as their line width and inversely as their length at ~ 1 pH/sq. Smaller metal line pitch will enable narrower, shorter inductors, with a smaller footprint. So reducing metal line pitch is an important factor.
- Adding more superconducting wiring layers. More wiring layers allow vertical structures to replace planar structures, increasing areal density. It will enable the vertical separation of power distribution and data transmission lines from the gates and free-up space for more gates, as well as reduce undesirable interactions between the gates and power/data signals. Smaller metal line pitch will also enable narrower, higher impedance transmission lines.

Bias Currents

A major cause of degraded margins in large circuits is inductive coupling of bias currents to gate inductors. Bias current increases approximately linearly with the number of JJs on a chip, at about 0.1 mA/JJ. For 10^3 junctions, the total current required is ~ 100 mA, at 10^4 JJ, ~ 1 A, and at 10^6 JJ, 100 A. Even small currents can affect shifts and reductions in margins in two ways: by direct inductive coupling to circuit inductors and by indirect coupling of uncontrolled ground return currents. These effects have frequently been ignored for small circuits, where they are small and are frequently compensated for by tweaking. By extension, they have frequently been neglected for large circuits. From discussions at the 2004 Applied Superconductivity Conference digital circuits session (Hypres, ISTEC-SRL, NG, and Chalmers), it appears that a major cause of degraded margins is due to bias currents.

- Direct coupling

Designers have implicitly assumed that mutual coupling between low impedance superconducting microstriplines is negligible. However, as total bias currents have increased and metal line pitch has decreased, this assumption has led to circuit failures. Northrop Grumman exhaustively tested several chips of a 4-bit slice of a MAC circuit (~ 1000 JJs) at low speed (ASC 2004). (Note that although testing was at low speed, all internal signals propagated at a high speed of ~ 50 ps/stage). Margins in one particular path through the circuit were reduced to about $\pm 6\%$; for other identical paths, they were comparable to design values of about $\pm 25\%$. Northrop Grumman was able to pinpoint the problem since the only difference was the position of nearby bias lines in the low margin path. At $1\text{-}\mu\text{m}$ spacing, the mutual inductance is $> 8 \times 10^{-15}$ H/ μm . Thus, 100 mA bias current coupling to a $10\text{ }\mu\text{m}$ line will produce a flux of approximately four magnetic flux quanta. Northrop Grumman concluded that the bias current shifted operating points and consequently reduced the margins.

The solution is to either place the bias lines far from circuit inductors or magnetically shield them. Accommodating large spaces is contrary to achieving dense chips. Shielding more than doubles the separation. One favorable solution is to locate current buses on separate layers, shielded from the active circuitry by a superconducting ground plane. Subterranean power lines (under the ground plane) isolated from the circuits have been proposed by several groups over the years; it has not yet been implemented in any process.

- Indirect coupling

A second method by which bias currents couple to gates is through return currents in the ground plane. It is common practice for SFQ chips to be installed in a multi-lead probe with a common ground return for all leads. One or more of these leads is used to supply the chip with power; others are used for input and output signals. Even when two leads are assigned for power, in and out, the power supply generally shares a common ground with other components. Return current distributes itself to all available ground leads to minimize the resistance. The problem is that on-chip, ground current migrates by many paths to minimize the inductance. These currents can couple into circuit inductors, shifting operating points, and consequently margins. This was conclusively observed by Lincoln Laboratory during the NSA Xbar project. The return current is not (as is frequently assumed) restricted to the ground plane either directly underneath the current lead, or to the edge of the chip. This effect was observed at Northrop Grumman for a modest size chip and was circumvented by tweaking bias currents (both positive and negative) down several lines until full circuit operation was achieved at speed. Of course, the latter method is not a practical solution. The solution is to separate the power supply from all other grounds (a floating supply) and place the bias + and – pins immediately adjacent to each other. In fact, one should make the bias supply leads coaxial, coplanar, or at least twisted to minimize stray magnetic fields from large current leads coming to the chip, and, use a coplanar or pseudo-coaxial pin arrangement. This was anticipated and implemented on FLUX-1.

Clocks, Jitter, and Timing

Clocks for SFQ circuits are themselves SFQ pulses, with similar power, pulse width, and jitter. Because high clock frequencies are required and the speed of transmission in stripline or microstrip is $\sim 1/3$ the speed of light in vacuum, clock signals cannot be supplied to an array of gates as they are in much slower CMOS circuits. Instead, they must be transmitted and distributed in a manner similar to data signals. Clock signals are best-generated on-chip using low-jitter circuitry, which have been demonstrated. Clock jitter and skew will be minimized if clock signals are generated on-chip, distributed on-chip, and distributed between chips similar to SFQ data transfer between chips.

Different sections of a large chip, as well as different chips, may have to operate asynchronously at 50 to 100 GHz clock frequencies and above. Clocks will have to be resynchronized at specific points in the circuit. Efficient methods to resynchronize and re-clock SFQ signals have been demonstrated.

Clock frequencies of 50 to 100 GHz will be limited more by circuit microarchitecture, timing and jitter in timing, and inter-gate delays, than by intra-gate delay. The inter-gate delay can be reduced by smaller gates placed closer together. This will depend on both fabrication technology and size-efficient gate design. Timing jitter is reduced by ballistic SFQ signal propagation between gates via matched passive transmission lines (PTL), rather than by active transmission lines consisting of a string of SFQ pulse repeaters (referred to as a Josephson transmission line, JTL). Since the junction impedance increases proportional to J_C , higher J_C enables narrower stripline/microstripline at the same dielectric thickness, again increasing potential circuit density. Most improvement in both gate density and clock frequency can be achieved by adding superconducting wiring layers as discussed above.

Jitter and timing errors are probably the most insidious factors in reducing margins as the clock frequency increases. Jitter occurs in all active devices: gates, fan-out and fan-in, and the clock distribution network. It is caused by all noise sources that can modulate the switching time: thermal noise in resistors, external noise, noise in the clock (particularly when externally supplied), and disturbs from clock and signal distribution. Jitter impacts the margins of clocked gates more than asynchronous gates because clocked gates need to synchronize the arrival of data and clock. It can reduce margins at high clock frequencies and therefore limit the useful clock frequency. Circuits are frequently designed without careful consideration of these jitter/noise sources. Consequently, when migrating to large circuits, margins could shrink rapidly.

Clock distribution is an important source of jitter if a large number of JTLs and splitters are used to propagate and distribute the clock. Every stage contributes jitter that accumulates as square-root of the number of stages. The use of PTLs instead of active JTLs will alleviate one source. However, splitters required to regenerate the clock remain a major source of jitter. A multi-line clock from a coherent clock array could reduce this jitter. Nevertheless, jitter will always be the ultimate limit on performance at high clock frequency.

Timing errors in design can be fatal and will occur if rigorous timing analysis is not part of the circuit design methodology. Commercial tools such as VHDL are available, but have not been widely used. Because of jitter, precise timing cannot be ensured. So, timing-error tolerant design should be used in critical circuits. Several methods to ensure proper data/clock timing within a clock cycle that add minimal circuit overhead have been demonstrated.

Other Noise Sources

In addition to ubiquitous thermal noise, there is noise in all input/output lines, including power lines, which feed back into the SFQ circuits. Noise measurements of circuits operating at 4 K almost universally show elevated noise levels, with typical effective noise temperatures of ~ 40 K. The sources for such noise are magnetic field noise and noise introduced by lines connected to warm noise sources. Even if testing is performed in a shielded environment, every wire, cable, etc., from RT to 4 K serves as an antenna that pipes signals and noise into the 4 K circuit. It is essentially impossible to provide a DC to infrared filter even in shielded rooms. Moreover, the terminations of RT test equipment generate at least 300 K wideband noise. This needs to be filtered at low temperature, preferably at 4 K.

One way to avoid some of the “antenna” noise is for all digital data lines to have a digital “optical isolator” that transmits only the desired bits and rejects all analog noise. Optical interconnects, under consideration for the wideband data links, could provide this added benefit. Even if optical interconnects are not used for the RT to 4 K data links, they should be considered as RFI isolators within the shielded cryostat environment at 300 K.

Power lines are particularly susceptible to transmitting various noise signals. Because of the high currents, filtering at low temperature is more difficult. One concept is to bring power from RT to ~ 40 K as RF, filter the RF with a narrow-band high temperature superconductor filter, and convert the RF to DC at 40 K. This has several advantages, including power transmission at high voltage and low current to reduce ohmic heating and noise filtering. From ~ 40 K, one can use zero-resistance high temperature superconductor current leads to 4 K.

Power and bias current

Power is dissipated continuously in the series resistors used to bias each junction. $P \sim 0.5I_C V_{DC}$ per junction, where V_{DC} is the DC voltage power bus. This static power dissipation dominates on-chip power dissipation. Ohmic dissipation can be reduced by using the smallest I_C consistent with the required BER, using designs with the fewest junctions (e.g., using passive rather than active transmission lines), and reducing the on-chip voltage supply as much as possible consistent with gate stability and preserving margins.

For a 2 mV bias supply, the static power dissipation is 100 nW/JJ. If one accepts ~10% reduction in gate margin (e.g., from 30% to 20%), one can reduce V_{DC} to $\sim 10^{-5}F(\text{GHz})$. At 50 GHz, $V_{DC} = 0.5$ mV and the static power dissipation is ~ 35 nW/JJ. Power is dissipated in every JJ in both the logic and clock networks at the SFQ switching rate, $P = 2 \times 10^{-15} f_{SFQ} I_C = 2 \times 10^{-19}$ watts/Hz per JJ (~ 1 electron-volt) or 0.2 nW/GHz/JJ. At 50 GHz, the irreducible SFQ power is $= 10$ nW/JJ.

A recent concept (called SCCL, Self-Clocked Complementary Logic) mimics CMOS by replacing bias resistors with junctions, forming a complementary pair wherein one and only one of the pair switches each clock cycle, and eliminates the ohmic loss. A few gates have been simulated and limited experiments have been performed successfully, but it has not been developed for a general gate family. An added value of this approach is incorporation of the local clock into the gate, eliminating the need for separate clock distribution circuitry. Normally, clock distribution has a splitter per gate that increases ohmic power, is an added source of jitter that reduces margins in large circuits, and occupies valuable space on-chip.

Chips with a large number of junctions require large bias currents. Even assuming all JJs are at minimum I_C of 100 μA , a 10^6 -JJ chip will require 100 A. Efficient methods of supplying bias current to the chips will have to be demonstrated for large circuits. A method to bias large circuit blocks in series (referred to as current recycling or current re-use) will be essential for large junction-count chips in order to reduce the total current supplied to the chip to a manageable value. Both capacitive and inductive methods of current recycling have been demonstrated at a small scale. Current recycling becomes easier at higher J_C . There have been no demonstrations of current recycling for large circuit blocks or for a large number of stages, so this will require development. Current recycling can reduce the heat load in the power lines into the cryostat, but does not reduce the on-chip power dissipation.

For current re-use, the ground planes under adjacent circuit blocks must be separated and subsequent blocks biased in series. It will also be necessary to isolate SFQ transients between adjacent blocks. This may be achieved by low pass filters, but will need to avoid power dissipation in the filters. Series inductance could provide high frequency isolation; the inductors could be damped by shunting with suitable resistance, such that there is no DC dissipation. Capacitive filtering may be problematic.

Efficient methods of supplying bias current to the chips need to be demonstrated for large circuits. The problem is to supply a large, low noise, stable current to the chip through the thermal-mechanical interface. Except for minimizing the number of JJs, power reduction does not reduce the current required to power the chip. This is discussed below.

Flux trapping

All Josephson circuits are sensitive to local magnetic fields. This sensitivity derives from the small size of the magnetic flux quantum; one flux quantum is equivalent to 2×10^7 gauss in a 1-cm^2 area. (The magnetic field of the earth is ~ 0.4 gauss.) JJ circuits are shielded from local magnetic fields, such as the earth's field, by high permeability shields. The field inside "good" shields can be as low as a few milligauss. Flux trapping occurs at unpredictable locations in superconducting films as they are cooled through their superconducting transition temperature, T_C . Flux trapped in ground planes can shift the operating point of SFQ circuits. One method for alleviating this effect is to intentionally trap the flux in holes (called moats) placed in the ground plane. One tries to place the moats where the trapped flux will not affect the circuit. There is no standard system for designing and locating moats. When flux trapping is suspected as the problem for a chip under test, the procedure is to warm the chip above T_C and cool it again, possibly after changing its orientation or position. If the behavior changes every time, the conclusion is that flux trapping is at fault. The assumption is that flux trapping is not reproducible.

Various moat protocols have been proposed and tried: long narrow channels, an array of circular holes, and random sizes and shapes. Studies at NEC and UC Berkeley concluded that long, narrow moats surrounding each gate are best, particularly if the area enclosed by the moat is smaller than the area of one flux quantum in the ambient field. However, moats cannot fully surround a gate because that would require wideband data lines to cross a region without a ground plane. Therefore, moats have breaks.

For a high probability of trapping flux in the moat, the moat inductance should be large to minimize the energy require to trap the flux. Thus, $E_{\text{SFQ}} = \Phi_0^2/2L_{\text{moat}}$ should be small. Narrow moats increase packing density, but have low inductance per unit length; long moats increase inductance. Therefore, long, narrow moats appear to be the best option. When flux is squeezed into the moat, it creates a corresponding shielding current in the surrounding superconducting film. Just as ground-plane currents are not confined either under a microstripline or to the edge of the ground plane, the shielding currents surrounding a moat are not confined to a London penetration depth from the edge of the moat. They are spatially distributed on the ground plane as required to meet the boundary condition that the normal component of magnetic field is zero. Assuming the length factor for shielding ground currents is approximately the width of the moat, the moat should be as narrow as litho permits and at least one width's distance removed from the nearest circuit inductor.

A second source of trapped flux is the equivalent of ESD in semiconductors. Transient signals that are not well filtered in any lead can produce trapped flux in the superconducting films and shift operating margins. The solution here is to employ isolation techniques in addition to filtering, including optical isolators in input/output lines and transmitting bias current to the shielded cryogenic environment by narrow band filtered RF.

Chips are usually flipped and bonded to a substrate with a superconducting film directly below the chip. This brings the active circuits at most a few microns from another superconducting ground plane. Although there is no data that supports this conjecture, one should be concerned and it needs to be addressed. EM simulations should help, but it will have to be resolved experimentally.

Appendix H

MRAM TECHNOLOGY FOR RSFQ HIGH-END COMPUTING

MRAM PROPERTIES AND PROSPECTS

Background

Magnetoresistive random access memory (MRAM) technology combines a spintronic device with standard silicon-based microelectronics to obtain a unique combination of attributes. The device is designed to have a large magnetoresistance effect, with its resistance depending on its magnetic state. Typical MRAM cells have two stable magnetic states. Cells can be written to a high or low resistance state and retain that state without any applied power. Cells are read by measuring the resistance to determine if the state is high or low. This resistance-based approach is distinctly different from common commercial memories such as DRAM and flash, which are based on stored charge. MRAM has made steady improvement and attracted increasing interest in the past ten years, driven mainly by improvements in thin-film magnetoresistive devices. There are two main approaches to MRAM: MTJ MRAM and SMT MRAM discussed below.

MTJ MRAM

The MRAM technology closest to production at this time has one magnetic tunnel junction (MTJ) and one read transistor per memory cell. As shown in Figure 1, the MTJ is sandwiched between two ferromagnetic films, whose relative polarization states determine the value of the bit. If the two polarizations are aligned, highly spin-polarized electrons can more easily tunnel between the electrodes, and the resistance is low. If the polarizations are not aligned, tunneling is reduced and the resistance state is high. MTJ materials in such circuits typically have aluminum oxide tunnel barriers and have a tunneling magnetoresistance ratio (TMR, the ratio of the resistance change to the resistance of the low state) of 25% to 50%. Recent demonstrations of TMR > 200% using MgO-based material indicate that large improvements in signal may be possible, thereby READ performance and scaling of future MRAM technology. The detailed scheme for programming the memory state varies, but always involves passing current through nearby write lines to generate a magnetic field that can switch the magnetic state of the desired bit. A recently improved MRAM cell, employing a synthetic antiferromagnet free layer (SAF), has demonstrated improved WRITE operation. A 4Mb MRAM circuit, employing a toggle-write scheme with a SAF layer, has shown improved data retention and immunity to cross-talk during the WRITE operation in a dense cross-point write architecture. The SAF layer also is expected to enable scaling of this MRAM architecture for several IC lithography generations.




Figure 1. Schematic of a 1T-1MTJ MRAM cell structure showing the sense path and programming lines. This cell has one MTJ and one transistor for read, and two orthogonal programming lines that are energized from transistors at the edge of the memory array.

SMT MRAM

Another direct selection scheme makes use of the recently observed interaction between a spin-polarized current and the bit magnetization, through so-called spin-momentum transfer (SMT). If a spin-polarized current is incident on a magnetic bit, the spin of the current-carrying spin-polarized electrons will experience a torque trying to align the electrons with their new magnetic environment inside the bit. By conservation of angular momentum, the spin-polarized current will also exert a torque on the bit magnetization. Above a critical current density, typically 10^7 to 10^8 A/cm², the current can switch the magnetic polarization by spin-transfer. SMT switching is most effective for bit diameters less than ~ 300 nm and gains importance as IC geometries continue to shrink.




Figure 2. Schematic of a proposed SMT MRAM cell structure showing the common current path for sense and program operations. Successful development of this architecture would result in a high-density, low power MRAM.

The main advantages of SMT switching are lower switching currents, improved write selectivity, and bits highly stable against thermal agitation and external field disturbances. However, SMT research is currently focused on understanding and controlling the fundamentals of spin-transfer switching. Several significant challenges remain before defining a product based on this approach. One such challenge is making a cell with a reasonable output signal level, as observations of spin-transfer so far have been confined to GMR systems, which have lower resistance, MR, and operating voltages compared to magnetic tunnel junctions. However, MTJ materials have very recently been identified with $MR > 100\%$ at very low resistances, and $MR > 200\%$ with moderate resistances. Such materials could potentially improve both the signal from SMT devices and lower the minimum currents needed for switching. Ongoing R&D in this area is directed at further decreasing the minimum switching currents and establishing materials with the necessary reproducibility and reliability for the low-resistance range. The outcome of these R&D activities will determine which SMT architectures can be used and at what lithography node. If switching currents are driven low enough for a minimum-size transistor to supply the switching current, this technology will meet or exceed the density of the established semiconductor memories, while providing high speed and non-volatility. Although SMT technology is less developed than MTJ MRAM is, it has the potential for higher density and orders-of-magnitude reduction in power dissipation.

We note here that for the specific case of MRAM integrated with RSFQ electronics, the low resistance of the all-metal GMR devices (instead of the MTJ) may be beneficial, as described in the follow-on section Integrated MRAM and RSFQ Electronics.

Potential for Scaling and Associated Challenges

Figure 3 compares the estimated cell size for standard (1T-1MTJ) MRAM, SMT MRAM using a single minimum-sized read/program transistor per cell, and NOR Flash. SMT devices have been demonstrated, but integrated SMT MRAM circuits have not, so that curve assumes technical progress that enables this dense architecture. The solid lines for the other two technologies indicate the goals and the dashed lines indicate how the cell size might scale if current known challenges are not completely overcome.




Figure 3. Estimated cell sizes for the two MRAM technologies compared to NOR Flash for IC technology nodes from 90nm to 32nm. Dashed lines indicate possible limitations to scaling if some materials and magnetic process challenges remain unresolved.

Scaling MRAM to these future technology generations requires continued improvement in controlling the micro-magnetics of these small bits and the MTJ material quality. The challenges are more difficult for SMT MRAM due to the smaller bit size required to take advantage of the smaller cell, and the need for increased SMT efficiency to enable switching at lower current densities. These differences will put more stringent requirements on the intrinsic reliability and on the quality of the tunnel barrier for SMT devices. Since the SMT technology is less mature than standard MRAM, it also will be necessary to consider multiple device concepts, material stacks, and corresponding SMT-MRAM architectures, so that optimal solutions can be identified.

Speed and Density

The first planned product from Freescale Semiconductor is very similar to their 4Mb Toggle MRAM demonstration circuit, which had a $1.55\text{-}\mu\text{m}^2$ cell size in a $0.18\text{ }\mu\text{m}$ CMOS process. This circuit had symmetric read and write cycle times of 25 ns. The Freescale circuit was designed as a general-purpose memory, but alternate architectures can be defined to optimize for lower power, higher speed, or higher density. Each optimization involves engineering trade-offs that require compromise of the less-important attributes. For example, a two-MTJ cell can be used to provide a larger signal for much higher speeds, but it will occupy nearly double the area per cell. A 128 kb high-speed circuit with 6-ns access time has recently been demonstrated by IBM. Since much of the raw signal (resistance change) is consumed by parametric distributions and process variations, relatively small increases in MR will result in large improvements in read speed. MTJ materials with four times higher MR have recently been demonstrated. It is therefore reasonable to expect continued improvements in access time especially at cryogenic temperatures where bit-to-bit fluctuations are dramatically reduced. Table 1 compares the performance of several memory technologies at the 90-nm IC node at room temperature, including that for a general-purpose MRAM and a high-speed MRAM. (Stand-alone 90-nm CMOS memory is in production today, although the embedded configuration is yet to be released.) Note that DRAM and Flash cell sizes are dramatically larger when embedded, as compared to the stand-alone versions. On the other hand, due to their backend integration approach, MRAM cell sizes remain the same for embedded and stand-alone.

Table 1. Comparison of MRAM with semiconductor memories. The MRAM column represents a general-purpose memory, while the High-Speed MRAM column represents architectures that trade some density for more speed. Either type of MRAM can be embedded or stand alone. Stand-alone Flash and DRAM processes become more specialized in order to achieve high density, and have much lower density when embedded, due to differences in fabrication processes compared to the underlying logic process. The density and cell size ranges for these two latter memories are very large due to the compromise needed for embedded fabrication.

	MRAM	MRAM	High-Speed MRAM	FLASH	SRAM	DRAM
Technology Node	0.18 μm Demo	90 nm Target	90 nm Target	90 nm Typical	90 nm Typical	90 nm Typical
Density (Mb)	1 - 32	16 - 256	4 - 32	4* - 4000	4 - 64	16* - 1,000
Wafer Size (mm)	200	200/300	200/300	200/300	200/300	200/300
Cycle Time (ns)	5 - 35	5 - 35	1 - 6	40-80 (Read) $\sim 10^6$ (Write)	0.5 - 5	6 - 50
Array Efficiency	40% - 60%	40% - 60%	40%-60%	25% - 40%	50% - 80%	40%
Voltage	3.3V/1.8V	2.5V/1.2V	2.5V/1.2V	2.5V/1.2V 9V - 12V internal	2.5V/1.2V	2.5V/1.2V
Cell Size (μm^2)	0.7 - 1.5	0.12 - 0.25	0.25 - 0.50	0.1 - 0.25*	1 - 1.3	0.065-0.25*
Endurance (cycles)	$>10^{15}$	$>10^{15}$	$>10^{15}$	$>10^{15}$ read, $<10^6$ write	$>10^{15}$	$>10^{15}$
Non-Volatile	YES	YES	YES	YES	NO	NO

*marks the embedded end of the range.

Integrated MRAM and RSFQ Electronics

MRAM circuits are presently fabricated in a back-end process, after standard CMOS electronics processing. This integration scheme is one of the reasons that MRAM is viewed as a strong candidate for a future universal embedded memory. MRAM could potentially be fabricated as a front-end process to a superconducting back-end technology (or vice versa), providing a high-performance, nonvolatile, monolithic memory solution for RSFQ technology. Some of the issues that must be addressed in pursuing this monolithic option are: effects of low temperatures on MRAM, compatibility of the required currents, power dissipation, and resistance matching. In addition, one needs to consider compatibility of MRAM and RSFQ processing with respect to materials and processing temperatures.

MRAM devices have higher MR at low temperature due to a reduction of thermal effects that depolarize the spin-polarized tunneling electrons. A 50% increase in MR from RT to 4.2 K is typical, e.g., from MR = 40% to MR > 60%. The higher MR combined with lower noise levels inherent at low-temperature operation would provide a much larger useable signal, and therefore much faster read operations. Temperature also has a big effect on the magnetics of the MRAM devices. Because of their small magnetic volume, these devices experience significant thermal fluctuations at room temperature, which increase the requirements for minimum switching field and minimum layer thickness to prevent thermally-activated write errors. MRAM bits for cryogenic temperatures can be designed with thinner layers and lower switching fields, reducing write currents. In large arrays, one always observes a distribution of switching fields. These distributions require significantly higher write currents than the mean switching current, typically at least 6σ above the mean for Mbit memories. The two main sources of these distributions are thermal fluctuations and micromagnetic variations, because bits in arrays are not exactly identical. At cryogenic temperatures, the thermal contribution will be negligible, leaving only the contribution from micromagnetic variation. Thus, unless there are unforeseen micromagnetic issues, the magnetic distributions should be narrower at low temperatures, leading to a further reduction in the write current. Overall, one might expect a 30% to 50% reduction in write current at cryogenic temperatures for standard MRAM and 10% to 30% for SMT MRAM.

Figure 4 shows how the minimum write-current for an SMT MRAM element would scale with IC lithography in a one-transistor-per-cell architecture. This plot assumes some improvement in spin transfer efficiency in order to complement the resistance of a minimum-sized pass transistor. The SMT cell is designed to maintain an energy barrier of 90 kT at normal operating temperatures. Thus, a $2\text{-}\Omega$ SMT device, 90 nm in size, would require a bias of 0.2 mV to supply the desired switching current of 0.10 mA. The same resistivity material at the 45-nm node would have a device resistance of $8\text{ }\Omega$; the corresponding switching current of ~ 0.04 mA would require a bias of only 0.24 mV. Since these are small bias values, it may be safe to conclude that providing the necessary drive currents for SMT cells from RSFQ circuits is less of an issue than for SMT MRAM-on-CMOS, where larger voltages are required to overcome ohmic losses. Moreover, as described before, at lower temperatures, the energy barrier requirement becomes essentially non-existent, so that the free layer can be made much thinner, limited only by the practical thickness needed for a good quality film. With further reduction in device dimensions below 90 nm, the thinner free layer would decrease the magnetic volume by \sim another factor of two, leading to a reduction in the critical current of up to 50% at cryogenic temperatures, compared to the numbers on this plot. Thus, it should be possible to directly connect superconducting JJ electronics to SMT MRAMs, even without the improvements in spin-transfer efficiency desired for room-temperature SMT operation.




Figure 4. Estimated switching currents for SMT devices at various lithography nodes. The device is designed to maintain a fixed energy barrier, as shown by the red line (right scale), for sufficient stability against thermally-activated errors at operating temperatures above room temperature.

Resistance matching and compatibility of the MRAM write current requirements to RSFQ electronics are topics of research for developing embedded MRAM-in-RSFQ circuits. Since high-density MRAM always employs the current-perpendicular-to-plane (CPP) device geometry, the resistance of a device scales as the inverse of the device area. Typically the material is characterized by its resistance-area product (RA), so that the device resistance (R) is given by $R=RA/\text{area}$. Since the resistance of the pass transistor used for MRAM-on-CMOS is in the $k\Omega$ range, it is desirable to have R in the $k\Omega$ or tens of $k\Omega$ range. This leads to a requirement for RA of several $k\Omega\text{-}\mu\text{m}^2$ for the current generation of MRAM, and scaling lower with subsequent technology generations. Very high-quality MTJ material can be made for this RA range, enabling progress of MRAM-on-CMOS. At the same time MTJ material for hard disk drive (HDD) read heads has been developed for a much lower RA, in the $< 10\text{-}\Omega\text{-}\mu\text{m}^2$ range, to meet the requirement for a $100\text{-}\Omega$ sensor in the head. Products with such heads have recently begun shipping, indicating some level of maturity in this type of material. While the requirements for the MTJ material used in HDD sensors are significantly different from MRAM requirements, these developments indicate that MTJ material development for lower RA is progressing at a rapid pace.

CPP-GMR material is not of practical use for MRAM-on-CMOS because the material is metallic, and therefore has very low $RA < 1\text{ }\Omega\text{-}\mu\text{m}^2$. In addition, typical MR values are $\sim 10\%$, compared to $\sim 50\%$ for standard MTJ material at room temperature. However, such material would easily provide a bit resistance on the order of 10Ω at advanced lithography nodes, providing a natural match for RSFQ circuitry. The lower MR may be acceptable as long as the bit-to-bit resistance uniformity is superior to that for MTJ bits, and given the low thermal noise available at cryogenic temperatures. It is not unreasonable to expect that the resistance distributions of GMR bits would be narrower than that for MTJ bits since the tunneling resistance depends exponentially on the local barrier thickness, while the resistance of the Cu barrier used in GMR material is a small part of the device resistance. In addition, defects in the tunnel barrier can cause dramatically lower resistance, while defects in the metal layers of a GMR device make only minor contributions. Of course, several other criteria must be met before a definitive choice between the MTJ and GMR approaches can be made. However, from these basic arguments it is apparent that GMR materials should be considered seriously for MRAM-in-RSFQ circuits.

Stand-alone Cryogenic MRAM

Whether the MRAM is embedded in a superconductive logic circuit or in a stand-alone RSFQ-MRAM chip, the general considerations for designing MRAM for low-temperature operation are identical to those outlined in the previous section. The ideal way to approach this would be to start with the best MRAM technology available at the time, MTJ or SMT-based, and modify it for lower power operation at low temperatures. At a minimum, a custom circuit design and optimized devices would be needed, as well as a program for performance characterization at the temperatures of interest.

The use of stand-alone memory would require a high-speed interface/bus between the RSFQ processors and the memory.

We acknowledge the contribution of Jon Slaughter, Freescale Semiconductors, Inc. in preparing the tutorial on MRM.

Appendix I

Superconductor Integrated Circuit Fabrication Technology

LYNN A. ABELSON AND GEORGE L. KERBER, MEMBER, IEEE

Invited Paper

Today's superconductor integrated circuit processes are capable of fabricating large digital logic chips with more than 10^8 K gates/cm². Recent advances in process technology have come from a variety of industrial foundries and university research efforts. These advances in processing have reduced critical current spreads and increased circuit speed, density, and yield. On-chip clock speeds of 60 GHz for complex digital logic and 750 GHz for a static divider (toggle flip-flop) have been demonstrated. Large digital logic circuits, with Josephson junction counts greater than 60 k, have also been fabricated using advanced foundry processes. Circuit yield is limited by defect density, not by parameter spreads. The present level of integration is limited largely by wiring and interconnect density and not by junction density. The addition of more wiring layers is key to the future development of this technology. We describe the process technologies and fabrication methodologies for digital superconductor integrated circuits and discuss the key developments required for the next generation of 100-GHz logic circuits.

Keywords—Anodization, critical current, flip-flop, foundry, interlevel dielectric, Josephson junction, niobium, niobium nitride, 100-GHz digital logic, photolithography, planarization, quantum computing, qubit, rapid single-flux quantum (RSFQ), reactive ion etch, resistor, SiO₂, superconductor integrated circuit, trilayer.

I. INTRODUCTION

In the past ten years, low-temperature superconductor (LTS) integrated circuit fabrication has achieved a high level of complexity and maturity, driven in part by the promise of ultrahigh speed and ultralow power digital logic circuits. The typical superconductor integrated circuit has one Josephson junction layer, three or four metal layers, three or four dielectric layers, one or more resistor layers, and a minimum feature size of 1 μ m. Niobium, whose transition temperature is 9 K, has been the preferred superconductor

due to its stable material and electrical properties, and ease of thin-film processing. The Josephson junction, which is the active device or switch, consists of two superconducting electrodes (niobium) separated by a thin (~ 1 nm thick) tunneling barrier (aluminum oxide). Josephson junctions, fabricated in niobium technology, exhibit remarkable electrical quality and stability. Although the physical structure of the Josephson junction is simple, advanced fabrication techniques have been developed to realize a high level of integration, electrical uniformity, and low defects. Today, niobium-based VLSI superconductor digital logic circuits operating at 100 GHz are a near-term reality and could have a significant impact on the performance of future electronic systems and instrumentation if the rate of innovation and progress in advanced fabrication continues at a rapid pace.

The promise of ultrahigh speed and ultralow power superconductor digital logic began in the mid-1970s with the development of Josephson junction-based single-flux quantum (SFQ) circuits. In 1991, Likharev and Semenov published the complete SFQ-based logic family that they called rapid SFQ (RSFQ) logic [1]. In 1999, researchers demonstrated a simple RSFQ T flip-flop frequency divider (divide-by-two) circuit operating above 750 GHz in niobium technology [2]. In terms of raw speed, RSFQ logic is the fastest digital technology in existence [3]. RSFQ logic gates operate at very low voltages on the order of 1 mV and require only about 1 μ W for even the fastest logic gates. As a result, on-chip RSFQ logic gate density can be very high even for 100-GHz clock frequencies. RSFQ logic is considered to be the prime candidate for the core logic in the next generation of high-performance computers [4]–[7] and is recognized as an emerging technology in the Semiconductor Industries Association roadmap for CMOS [8]. Recently a prototype of an 8-b, RSFQ microprocessor, fabricated in the SRL (formerly NEC) 2.5-kA/cm² process (see Table 1), demonstrated full functionality at a clock frequency of 15.2 GHz with power consumption of 1.6 mW [9]. Even smaller scale applications, such as ultrawide band digital communication

Manuscript received December 3, 2003; revised April 16, 2004.

L. A. Abelson is with Northrop Grumman Space Technology, Redondo Beach, CA 90278 (e-mail: lynn.abelson@ngc.com).

G. L. Kerber was with Northrop Grumman Space Technology, Redondo Beach, CA 90278 USA. He is now in San Diego, CA 92117 USA (e-mail: george.kerber@glkinst.com).

Digital Object Identifier 10.1109/JPROC.2004.833652

Table 1
Representative Nb and NbN-Based IC Processes

Institution	No. Masks	Technology	J_c (kA/cm ²)	Min. JJ Area (μm^2)	Min. Feature (μm)	Wire Layers	Resistors	Ground-plane	Process optimized for ¹
AIST [26]	7	Nb (SIS)	1.6	7.8	1.5	2	1	Bottom	RSFQ logic
Hypres (1.0kA) [28]	10	Nb (SIS)	0.1, 1, 2.5	9.0	2.0	3	2	Bottom	QC, RSFQ logic, ADC's
Hypres (4.5kA) [29]	11	Nb (SIS)	4.5, 6.5	2.25	1.0	3	2	Bottom	RSFQ logic, ADC's
IPHT Jena [34]	12	Nb (SIS)	1	12.5	2.5	2	1	Bottom	RSFQ logic
Lincoln Lab [21]	8	Nb (SIS)	0.1, 0.5, 10	0.5	0.7	2	2	Top	OC, RSFQ logic
ISTEC (SRL) [33]	9	Nb (SIS)	2.5	4.0	1.0	2	1	Bottom	RSFQ logic
NGST [18]	14	Nb (SIS)	8	1.2	1.0	3	2	Bottom	RSFQ logic, ADC's
NGST	12	Nb (SIS)	0.1	0.8	1.0	3	1	Bottom	QC
SUNYSB [32]	8	Nb (SIS)	0.2 to 12	0.06	0.25	2	2	Top	QC, R&D
NIST [35]	8	Nb (SNS)	200	1.0	0.7	2	1	None	Voltage standard
NIST [20]	10	Nb (SIS)	0.5	2.0	1.0	2	1	None	Squids
PTB [23], [24]	8	Nb (SINIS)	1	12	2.5	1	1	Bottom	RSFQ logic
PTB [22]	8	Nb (SIS)	1	12	2.5	1	1	Bottom	RSFQ logic
UC Berkeley [25]	10	Nb (SIS)	10	1.0	1.0	2	1	Bottom	RSFQ logic
Univ. Karlsruhe [36]	7 or 8	Nb (SIS)	1 to 4	4.0	1.0	2	1	Bottom	Analog, RSFQ logic
NGST [31]	12	NbN (SIS)	1	7.1	2.0	3	2	Bottom	RSFQ logic, ADC's
CRL [37]	8	NbN (SIS)	2.5	4.0	2.0	2	1	Bottom	RSFQ logic
AIST [27]	7	NbN (SNS)	30	16	1.5	1	1	None	Programmable voltage standard
CAE [30]	10	NbN (SIS)	5	2.5	2.0	3	2	Bottom	RSFQ logic

¹ QC = quantum computing

systems and ultrafast digital switching networks will benefit from its unparalleled speed and low power that far outweighs the need to provide 4 or 10 K operating environment [10], [11]. RSFQ logic may play a key role in future quantum computers as readout and control circuits for Josephson junction-based qubits [12]–[14].

For these reasons, RSFQ logic has gained wide acceptance. In the past ten years, many organizations have developed and sustain advanced fabrication processes or are developing their next-generation process tailored to RSFQ logic chip production (see Table 1). Both niobium and niobium nitride superconductor material technologies are supported. Several of the organizations in Table 1 provide valuable foundry services to the industrial and academic communities. For example, HYPRES, in the United States, has tailored its 1-kA/cm² fabrication process to allow a large variety of different chip designs on a single 150-mm wafer. Although the fabrication technology is not the most advanced, the cost per chip is low, which is particularly attractive to the academic community for testing a new idea or design. HYPRES is also developing an advanced 4.5-kA/cm² foundry process and a very low current density, 30-A/cm² process for the development of quantum computing circuits. Foundry services, available from SRL, are enabling many groups in Japan to design and test new circuit concepts. The SRL foundry offers a 2.5-kA/cm² process and is developing its next-generation 10-kA/cm² process.

The full potential of digital superconductor logic-based systems can only be realized with advanced chip fabrication technology [15]. In the United States, advances in chip fabrication are being driven in part by the need to demonstrate niobium-based digital VLSI logic chips for high-performance, petaFLOPS computing using the Hybrid Technology Multithreaded (HTMT) architecture [16]. Under the aggressive schedule of the HTMT architecture study program, chip fabrication technology advanced by two generations in junction current density (J_c) from 2 kA/cm² to 8 kA/cm² [17]. A 12-stage static divider or counter, fabricated in the 8-kA/cm² process, demonstrated correct operation from ~dc to 300 GHz [18], and a more complex RSFQ logic circuit achieved 60-Gb/s operation at low bit error rates [19].

Advanced fabrication processes simultaneously have reduced junction size, kept junction critical current (I_c) spreads below 2% (1σ), and improved chip yields compared to a decade ago. These advances have come from cooperative efforts between industrial and university research groups. The next-generation 20-kA/cm² 0.8- μm junction, six-metal layer niobium process is expected to achieve on-chip clock rates of 80–100 GHz and gate density greater than 50 000 gates/cm². This paper describes the present, well-established niobium-based chip fabrication technology and the roadmap to a 20-kA/cm² 0.8- μm process. We also discuss niobium nitride-based chip fabrication, which is much less mature compared to




Fig. 1. Cross section of NGST's 8-kA/cm² niobium-based 14-mask step integrated circuit process, showing anodized ground plane (Nb₂O₅), junction anodization, three niobium wiring layers, bias-sputtered SiO₂ insulation, and two resistor layers (MoN_x and Mo/Al).

Table 2
Nb Process Layers

Layer	Material	Thickness
GND (ground)	Nb	150 nm
GNDC (1st ground insulation)	Nb ₂ O ₅	144 nm
RESI (resistor 1)	MoN _x	95 nm
RESL (resistor 2)	Mo/Al	65 nm
SIOG (2nd ground insulation)	SiO ₂	150 nm
Trilayer-Base Electrode	Nb	150 nm
Trilayer- Tunnel Barrier	Al/AlO _x	8 nm
Trilayer-Counter Electrode	Nb	100 nm
SIOA (1st Insulation)	SiO ₂	200 nm
WIRA (1st Wireup)	Nb	300 nm
SIOB (2nd Insulation)	SiO ₂	450 nm
WIRB (2nd Wireup)	Nb	600 nm
GOLD (Pad Contact)	Ti/Pd/Au	40/400/40 nm

niobium, but niobium nitride is important due to its higher (10 K) operating temperature.

II. LTS PROCESS TECHNOLOGIES

A. Niobium and Niobium Nitride Process: An Overview

LTS integrated circuit technologies are well established at a variety of industrial and research institutions. Table 1 summarizes the salient features of representative integrated circuit processes from around the world [18], [20]–[37]. Most of the processes are based on niobium technology, with a few groups actively pursuing development of niobium nitride-based process technology. Independent of whether the process is niobium or niobium nitride, these processes have converged to a similar topology. Fig. 1 shows a cross section of the niobium-based process at Northrop Grumman Space Technology (NGST) and Table 2 lists the typical layer characteristics, which are representative of the processes in Table 1. These processes have been used to demonstrate a wide range of analog and digital circuits operating at 4 and 10 K. The recent interest in quantum computing has lead

Table 3
Minimum Feature Sizes and Reticle Sizing

Mask Name	Min. Feature (Drawn) (μm)	Min. Space (Drawn) (μm)	Reticle Sizing ¹ (per side) (μm)
GNDE	2.0	2.0	-0.1
GNDC	2.5	2.5	+0.2
TRCH	2.0	2.0	-0.1
RESI	2.0	1.5	+0.0
RESL	3.0	2.0	+0.2
SIOG	1.5	1.0	+0.0
JUNC	1.2	1.5	+0.1
JNC2	N/A	N/A	+0.8
TRIW	1.3	1.3	+0.1
SIOA	1.0	1.0	+0.0
WIRA	1.3	1.3	+0.1
SIOB	1.5	1.5	+0.0
WIRB	2.0	2.0	+0.3
GOLD	4.0	3.0	+0.0

¹ Sizing limited to increments of ±0.1 μm due to e-beam spot size.

some institutions to adapt their 4 K, RSFQ processes for millikelvin operation. Due to the difficulty of heat removal at millikelvin temperatures and to the sensitivity of qubit decoherence from heat dissipation, the fabrication processes must be optimized to produce RSFQ circuits with even lower power dissipation (lower current density, low noise resistors, etc.) than the typical 4 K RSFQ circuit [13].

Process complexity is indicated both by the number of masking levels (from 7 to 14, depending on the process) and the minimum feature size (from 0.25 to 2.5 μm). Masking levels and feature sizes for the NGSF niobium process are summarized in Table 3. The specific details of any particular process depend on the type of circuit application and the availability of process tools. Most of the niobium processes use superconductor–insulator–superconductor (SIS) tunnel junctions, based on Nb/Al–AlO_x/Nb trilayers for the Josephson junction, although there is work on alternate devices based on superconductor–normal conductor–superconductor (SNS) or superconductor–insulator–normal conductor–insulator–superconductor (SINIS) where “N” stands for a normal conductor [27], [35], [38]–[40]. The niobium nitride processes use either MgO or AlN as the tunnel barrier material in the Josephson junction. In addition to the trilayer, these processes employ superconducting interconnect (one to three levels) of either niobium or niobium nitride wires, thin-film resistors composed of a variety of normal metals (e.g., Pt, Mo, NbN_x, MoN_x, TiPd), and interlevel dielectrics of either SiO or SiO₂. Anodized niobium (Nb₂O₅), due to its high dielectric constant, is used in some of the processes for forming capacitors and for isolation between wiring levels. Most of the processes have a separate niobium or niobium nitride ground plane layer for shielding magnetic fields and controlling circuit inductances. The

ground plane may be located either below the trilayer or, less often, on the top. Optical photolithography (either g-line or i-line) is used to transfer the mask design to the photoresist in most cases. E-beam lithography is also used to write the smallest features ($<1\ \mu\text{m}$) in some processes. The features are then patterned by a variety of methods including etching by reactive ion etching (RIE) or inductively coupled plasma (ICP) etching, anodization, and liftoff. The integrated circuit fabrication processes are discussed in more detail below.

B. Fabrication Challenges

The challenges in the fabrication of superconductor integrated circuits fall broadly into two categories: improving parametric performance and minimizing process-induced defects. Parametric performance involves the targeting of important device parameters and minimizing device variation, both on a local and global scale. We discuss this in more detail in Section V. Process defects include unwanted contamination, lack of integrity in wiring or dielectrics, etc., and are typically represented as defect density. Even for modest feature sizes (i.e., $\sim 1\ \mu\text{m}$), defect density is an important consideration and care must be taken to minimize contributions from the environment, while working to reduce contributions from the process itself. These challenges are similar to those faced by the semiconductor industry, because many of the tools, methods, and materials are similar. We have adopted solutions already developed by the semiconductor industry and adapted them to the specific needs of superconductor integrated circuit manufacturing. For example, class 100 to class 10 clean rooms, depending on the level of integration, and clean room process tools, such as load-locked vacuum systems and automated wafer handling, are used to minimize defects.

Much work has been done over the past decade to minimize process-induced defects. Process improvements to address dielectric integrity included the use of sputtered SiO_2 as a replacement for evaporated SiO , which helped to reduce pinhole density in the interlevel dielectric and improved step coverage of the overlying metal layer. Implementation of bias-sputtered SiO_2 was another improvement that further improved dielectric integrity and improved metal step coverage [41]. Another approach to improved step coverage is the use of electron cyclotron resonance (ECR) plasma-enhanced chemical vapor deposition (PECVD) SiO_2 , which provides a collateral benefit of improved junction characteristics [20], [42].

Another challenge in superconductor integrated circuit fabrication is addressing material properties limitations for materials such as niobium nitride. Unlike niobium in which the Nb/Al- AlO_x /Nb trilayer is relatively straightforward to produce with good uniformity, the tunnel barrier for niobium nitride is more challenging. Sputter-deposited MgO and AlN are the typical material choices for tunnel barriers [37], [43]. Controlling the thickness (on the order of 1 nm) and uniformity (to better than 0.01 nm!) of an ultrathin barrier such as MgO is difficult by conventional sputter deposition techniques. Accurate targeting of junction I_c becomes very difficult because J_c is an exponential

Table 4
Nb and Al Sputter Deposition Parameters

Film	Nb	Al
Target purity	99.95% Nb	99.999% Al
Gun current	2.5 A (constant current)	2.75 A (constant current)
Gun voltage	300 to 320 V	400 V
Target-to-Substrate	4.5 in	7.0 in
Substrate Bias	0 V	0 V
Ar pressure	5.0 mTorr	5.0 mTorr
Ar purity	99.9999%	99.9999%
Deposition rates	1.1 nm/s	0.4 nm/s
Substrate Temperature	$\sim 25^\circ\text{C}$	$\sim 25^\circ\text{C}$

function of the tunnel barrier thickness. *In situ* oxidation of deposited Mg is another method that has been explored with encouraging results [44]. The potential advantage of this approach is a uniform barrier thickness across the wafer and more controllable J_c targeting. Although robust to chemical damage during processing, niobium nitride poses other difficulties for integrated circuit fabrication because of its large penetration depth and columnar growth [45], [46]. In order to overcome this, layer thicknesses are increased, causing step coverage problems. Innovative circuit design can help mitigate the problem, but the real solution is planarization and migration to other materials such as NbTiN, which has a much lower penetration depth, and so it can be made thinner [47]. Planarization, as discussed in Section IV, has been used successfully in niobium-based technology and could readily be adapted to niobium nitride to mitigate some of the step coverage problems.

III. JUNCTION FABRICATION

A. Nb/Al- AlO_x /Nb Trilayer Deposition

Fabrication of large numbers of Josephson junctions with predictable and uniform electrical properties is the key first step in the development of an advanced superconductor integrated circuit process. Fabrication of high quality Josephson junctions starts from an *in situ* deposited trilayer of Nb/Al- AlO_x /Nb. The trilayer is patterned using standard lithographic and RIE processes to define the niobium base and counterelectrodes of the Josephson junction. This has been the preferred method since the first demonstration of Nb/Al- AlO_x /Nb trilayer process [48]. Many details of trilayer deposition processes and basic junction fabrication can be found elsewhere [49]–[52].

The trilayer deposition is performed in a dedicated process tool (sputter deposition system) designed for this process, which is standard practice in the industry. The process tool generally consists of a multigun, sputter deposition chamber, oxidation chamber, glow discharge chamber, and a load lock chamber to transfer wafers in and out of the process tool. The process tool should be capable of maintaining base pressures in the low 10^{-8} torr range. In the trilayer deposition process used at NGST, the niobium base electrode and aluminum tunnel barrier metals (Nb/Al) are sputter-deposited sequentially in the deposition chamber.




Fig. 2. Current density (J_c) versus oxidation pressure-time product for J_c of 2–20 kA/cm². The • symbols indicate NGST data, and rectangular boxes outline all data from other processes [58].

Typical niobium base electrode thickness is 150 nm, and the aluminum thickness is 8 nm. The wafer is transferred to the oxidation chamber to partially oxidize the exposed aluminum layer to form a thin (~ 1 nm thick) aluminum oxide (AlO_x) tunnel barrier and transferred back to the deposition chamber to deposit the niobium top or counterelectrode layer to complete the *in situ* formation of the Nb/Al- AlO_x /Nb trilayer. The niobium counterelectrode thickness is typically 100 nm. In order to produce junctions of high electrical quality, it is important to optimize the argon sputter gas pressure to produce near zero stress in both niobium base and counterelectrode films [53]–[56]. Small junctions and submicrometer-sized junctions are particularly sensitive to film stress, which tends to increase subgap leakage current and decrease I_c uniformity [57]. Table 4 lists the optimized NGST deposition parameters.

Since J_c varies exponentially with barrier thickness, junctions require precise control over oxidation pressure, time, and temperature. The NGST oxidation chamber uses a mass flow controller to provide constant flow of high purity oxygen (99.999%) and a capacitance manometer and variable throttle valve connected in a feedback loop to dynamically control pressure. The feedback loop controls pressure to better than ± 0.1 mtorr. Typical oxidation pressure and time for the NGST 8-kA/cm² process are 9.0 mtorr and 30 min. For current densities greater than about 10 kA/cm², a dilute 10% O_2 , 90% Ar mixture is used in order to maintain a more favorable pressure range for feedback control.

The dynamic oxidation process used at NGST has excellent stability over time and good run-to-run repeatability. Fig. 2 shows the J_c dependence on oxidation pressure-time product for the NGST process and for several processes [58]. All data are in good agreement up to a J_c of 20 kA/cm².

To minimize run-to-run variations in J_c and to achieve low subgap leakage, it is important to keep the wafer temperature




Fig. 3. Trend chart of current density (J_c). Mean = 8.3 ± 0.37 kA/cm² or $\pm 4.4\%$ (1σ). Target specification = 8.0 kA/cm² $\pm 10\%$ (1σ).

constant and near room temperature, or colder if possible, either by active temperature control or by passive cooling. Active temperature control during deposition and barrier oxidation is highly desirable, but it is often not practical due to limitations of existing deposition tools. In the NGST trilayer deposition tool, the wafer and carrier are clamped to a large heat sink using indium foil backing, which minimizes the temperature rise and reduces temperature gradients across the wafer. The temperature rise during deposition is limited to a few degrees above room temperature and remains nearly constant during oxidation. For the 8-kA/cm² process, this method of passive cooling is sufficient to keep the average, run-to-run variations in J_c below 5% (1σ) as shown by the trend chart for J_c in Fig. 3. J_c is calculated from the slope of a least-squares fit of the square root of junction I_c versus junction diameter for five junction sizes. Error bars are determined from measurements of five chips distributed across the wafer. Across-wafer J_c spreads are as low as 1.5% (1σ).

B. Josephson Junction Fabrication

To improve circuit speed and performance, each new-generation process is based on increased junction




Fig. 4. Junction fabrication process showing key features: junction anodization [Fig. 4(b)] is self-aligned to junction; anodization etch [Fig. 4(c)] requires high selectivity to niobium; anodization ring is contained entirely within the space between edges of junction and base electrode [Fig. 4(d)]. (a) Deposit Nb/Al-AlO_x/Nb trilayer and apply junction mask. (b) Etch counterelectrode and anodize. (c) Apply second junction mask and etch anodization. (d) Remove photoresist mask to complete junction fabrication.

J_c . This requires junction sizes to decrease in order to maintain approximately the same I_c range. Since I_c is proportional to the square of the junction diameter, good dimensional control becomes one of the critical issues that affect I_c targeting and I_c spreads. For example, at 8 kA/cm² and minimum junction I_c of 100 μ A, variations in junction diameter must be controlled to less than $\pm 0.06 \mu$ m if variations in I_c are to remain under $\pm 10\%$. In addition, small junctions are also more sensitive to perimeter effects. Improvements in photolithography and RIE processes (see Section IV) have kept pace with the demand for improved dimensional control with little additional investment in new process tools.

In the NGST process, the junction is defined on the niobium counterelectrode of the trilayer by a photoresist mask shown in Fig. 4(a). Next, the niobium counterelectrode is dry etched in SF₆ [see Fig. 4(b)]. Since SF₆ does not etch aluminum or Al₂O₃, the etch stops on the tunnel barrier protecting the niobium base electrode. The dry etch process is performed in a simple parallel plate RIE tool. SF₆ etch chemistry in this tool produces clean, residue-free features that have vertical edges and no undercutting [59]. Another popular niobium dry etch gas is CF₄ which has been shown to produce residue free, submicrometer features in niobium

Table 5
Reactive Ion Etch Parameters

Etch Process	Gases ¹	Pressure (mTorr)	Power (W)	Etch Rate (nm/min)
Nb	SF ₆	15	30	95
MoN _x	SF ₆	25	75	28
Nb ₂ O ₅	CHF ₃ + 5% O ₂	100	150	5
SiO ₂	CHF ₃ + 27% O ₂	80	190 / 55	43 / 12

¹ O₂ percentage of total flow.

using an ECR plasma etch tool [60]. Immediately after etching, the wafers are lightly anodized to passivate the junctions, and then the photoresist is stripped. The anodization process protects the perimeter of the junctions from chemical attack during the photoresist strip and subsequent processing steps.

Many junction anodization processes have been described in the literature [48], [61]–[63], but only “light” anodization, described first by Gurvitch [48] and recently refined by Meng [64], offers protection from process damage and is scalable to submicrometer dimensions. Postetch junction passivation using “light” anodization has been a key development in the 8-kA/cm² process to minimize junction damage from subsequent wet processing steps. For example, AZ300T photoresist stripper [65] and deionized water rinse in combination can attack or erode the exposed aluminum and very thin (~ 2 nm) AlO_x tunnel barrier along the edge or perimeter of the junction. This can increase subgap leakage and degrade I_c spreads. The wafer is typically anodized in a mixture of ammonium pentaborate, ethylene glycol, and deionized water to 15 V. At 15 V the aluminum barrier metal is anodized completely to Al₂O₃ (~ 15 nm thick), and the exposed niobium layer is partially converted to Nb₂O₅ (~ 22 nm thick) [18]. Nb₂O₅ and Al₂O₃ make good passivation layers because of their resistance to attack by standard process chemistries.

In the anodization process described by Meng [64], Nb/Al-AlO_x/Nb junctions are formed with a self-aligned annulus that is lightly anodized to form an insulating double layer of Al₂O₃ and Nb₂O₅ on the bottom of the annulus (base electrode) and a single layer of Nb₂O₅ on the side-walls (counterelectrode). The anodization layer passivates the junction and sidewalls of the annulus. For the 8-kA/cm² process, NGST developed a variation of Meng’s light anodization process that uses a second junction-masking step. In the NGST process, shown in Fig. 4(c), the bulk of the anodization layer is removed everywhere on the counterelectrode except for the region around the junction protected by the second junction photoresist mask. A combination of wet dips in a dilute HF-nitric acid mixture and buffered oxide etch (BOE) is used to remove Al₂O₃ layer and a dry etch in CHF₃ + 5% O₂ (see Table 5) is used to remove the Nb₂O₅ layer. This creates a self-aligned, anodization layer or passivation ring around the junction as shown in Fig. 4(d). Fig. 5 is a scanning electron microscope (SEM) photograph of a




Fig. 5. SEM photograph of a 1.0- μm junction and self-aligned anodization ring on base electrode.




Fig. 6. SEM photograph of a partially completed T-flip flop stage showing junction, anodization ring, and base electrode.

1- μm junction surrounded by a 0.8- μm wide anodization ring. The second junction mask is designed to confine the anodization layer to fit entirely within the minimum design rule space (0.8 μm) between the edge of junction counter-electrode and the edge of the base electrode [66]. Therefore, it has no impact on circuit density and enables scalability to higher density processes as better photolithography tools become available. Since Al_2O_3 is an excellent insulator and etch stop, junction contacts can be slightly larger than the junction, and the alignment of contact to junction can be relaxed without risk of shorting between the wiring layer and base electrode. As a result, junctions become the minimum definable feature [18]. The SEM photograph in Fig. 6 illustrates the use of the lightly anodized junction process in a simple RSFQ divider circuit that has completed fabrication through base electrode etch.

C. Electrical Characteristics and I_c Spreads

The typical I - V characteristics of a minimum size, 1.25- μm diameter junction fabricated in NGST's 8-kA/cm²




Fig. 7. Typical current-voltage characteristics of a 8-kA/cm² 1.25- μm diameter junction. $V_m = 27$ mV, $V_g = 2.84$ mV, $\Delta V_g = 0.08$ mV, and $I_c R_N = 1.46$ mV.




Fig. 8. Typical current-voltage characteristics of a series array of 100, 8-kA/cm², 1.25- μm diameter junctions. I_c nonuniformity $\sim 1.6\%(1\sigma)$.

niobium process are shown in Fig. 7. The subgap leakage (V_m) is greater than 25 mV, where V_m is defined as the product of $I_c \times$ resistance measured at 2 mV in the voltage range below 2Δ , the energy gap. The deviation in junction diameter from its drawn size, which is determined from measurement of I_c and J_c , is only $+0.02$ μm (larger), and is typical of the critical dimension (CD) control. Based on electrical measurements of junction size from a large number of wafers and five sites per wafer, typical within-wafer variations are on the order of ± 0.02 μm , well within the ± 0.06 - μm specification.

Fig. 8 is a plot of the I - V characteristics of a series array of on hundred 1.25- μm diameter junctions with J_c of 8 kA/cm². The spread in I_c of the 100 junctions is 1.6% (1σ). The array occupies an area of about 200 $\mu\text{m} \times 200$ μm and has on-chip RC filters to suppress sympathetic switching [67]. Table 6 compares the best I_c spreads obtained in the 8-kA/cm² process to 2-kA/cm² and 4-kA/cm² processes, which did not use light anodization or improved low-power etch recipe [18]. Previous data suggested that I_c spreads increase rapidly for small junctions [67]; however, the 8-kA/cm² results clearly show that advanced processing can keep I_c spreads low.

Table 6
Junction Array Summary of Three NGST Process Generations

J_c (kA/cm ²)	2	4	8
Junction Diameter	2.50 μ m	1.75 μ m	1.25 μ m
No. Junctions	100	100	100
ΔI_c (1 σ) %	1.1	1.4	1.4
ΔI_c (max-min) %	± 2.5	± 3.4	± 3.5

IV. PROCESS INTEGRATION

A. Design Rules

Design rules provide details on circuit layout, minimum feature sizes, and electrical rules for the processes and are available from many institutions listed in Table 1. The importance of the design rule document is that it describes the process capability and provides the range of parameters expected on any given chip, given adherence to the guidelines specified. The NGST 8-kA/cm² niobium foundry process has demonstrated functional circuit yield at a level of integration of 2000–3000 junctions per chip and clock speeds of 300 GHz. The near-term goal was to yield circuits with greater than 60 000 junctions per chip utilizing existing fabrication tools. To reach this goal requires a very stable fabrication process, good CD control, predictable electrical performance, highly optimized design rules, and low defect density.

The NGST 8-kA/cm² process has 14 masking steps, which includes one ground plane, two resistor layers, and three wiring layers [66]. The minimum wire pitch is 2.6 μ m, and the minimum junction and contact sizes are 1.25 and 1.0 μ m, respectively. This process has demonstrated fabrication of junctions as small as 1.0 μ m and wire pitch of 2.0 μ m, but design rule minimum feature sizes were conservative to guarantee a high yield. Minimum feature size design rules are summarized in Table 3.

Design rules should be process-bias independent, i.e., drawn features are equal to final, on-wafer features. Process bias is the loss (or gain) in feature size between drawn and final on-wafer dimensions due to lithography and etch processes. The process bias is compensated for by sizing the reticle (mask). Reticle sizing for each layer is determined from a combination of electrical, optical, and SEM measurements. When process bias becomes a substantial fraction of the minimum feature size, further reductions in design rules are impossible even if the process tools are capable of defining smaller features. Therefore, lithography and etch processes should be optimized to produce minimum loss (or gain) in CD.

B. CD Control

NGST used a photoresist optimized for both g-line and i-line wavelengths [65]. Using this resist and g-line 1X Ultra-*tech* steppers, a resolution of 0.65 μ m (light field reticle) was achieved, and 1.0- μ m features could routinely be defined in




Fig. 9. SEM photograph showing the reentrant step coverage of sputter-deposited SiO₂ over a 500-nm metal step.

photoresist with minimum CD loss and excellent reticle linearity [59]. This photoresist is more than adequate for the 8-kA/cm² generation. Batch develop of the photoresist has been the standard practice, but spray develop has the potential to further reduce across-wafer CD variation and should become standard practice in the future.

All niobium layers, including junction and counterelectrodes, are reactively ion etched in SF₆. This dry etch process has been highly optimized to achieve an across-wafer (for 100-mm-diameter wafers) etch rate nonuniformity of less than 1%. The RIE tool operates at 30-W RF power and 15-mtorr pressure. Under these conditions, the RIE process produces little damage to the photoresist, and etched features (junctions and wires) have vertical sidewalls. The CD loss for the first and second wiring layers is less than 0.1 μ m and between 0.1 and 0.2 μ m for the third thicker wiring layer. Etch parameters for niobium and for other materials are summarized in Table 5.

CD control of contacts in the SiO₂ interlevel dielectric layers is not as critical. The minimum contact feature is 1.0 μ m, and the contact etch must simply clear a 1.0- μ m minimum opening. Contacts are etched in a mixture of CHF₃ + 27% O₂ to produce sloped walls and improve step coverage.

C. Interlevel Dielectric

The most common interlevel dielectric material for superconductor integrated circuit fabrication is sputter-deposited SiO₂ [68]–[70]. It has low defect density and can be deposited at low temperatures, since temperatures above about 150 °C can degrade the electrical properties of Nb/Al-AlO_x/Nb junctions. However, the step coverage of sputter-deposited SiO₂ is poor and would limit yield of an LSI or VLSI circuit process [41], [68]. The cross section of sputter-deposited SiO₂ over a step is shown in Fig. 9 and illustrates the “reentrant” structure that leads to poor step coverage by the next metal layer. Step coverage improves upon applying RF bias to the wafer during sputter deposition [71]. RF bias produces positive ion bombardment of the wafer surface and resputtering of the deposited oxide. The




Fig. 10. SEM photograph showing the improvement in step coverage of SiO_2 over a 500-nm metal step with substrate bias applied during deposition.

removal rate of SiO_2 at the surface of the wafer, relative to deposition, is an increasing function of RF bias. Within a certain range of deposition and removal rates and features sizes, bias-sputtered SiO_2 is locally planarizing. Local planarization of surface topology using bias sputtering has been studied extensively [72]–[74] and is due to the fact that under ion bombardment, sloped features are resputtered at a higher rate than flat areas because of the angular dependence of sputter yield which is maximum at about 65° for SiO_2 [75]. Bias-sputtered SiO_2 films also show improved microstructure, increased dielectric strength, reduced surface roughness, and reduced defect density [76], [77].

In the NGST process, a low-frequency (40-kHz) power source was used to supply bias to the wafer [69]. The dc self-bias of the wafer is used to monitor the deposition process, with feedback to the low frequency source to control power and, hence, removal rate. Atomic force microscope measurements indicate that the surface of the low-frequency bias-sputtered SiO_2 films are essentially featureless and have a surface roughness of less than 0.1 nm (rms). In contrast, the surface roughness of unbiased, sputtered SiO_2 films, deposited under same conditions and thickness, is on the order of 1.3 nm (rms). Additional details of the deposition tool, sputtering parameters, and film properties can be found in [69].

The improvement in edge profile of SiO_2 over a step using 40-kHz bias is shown in Fig. 10. Bias sputtering completely eliminates the reentrant oxide step, improves wire critical currents, and greatly reduces the probability of metal bridging over steps and shorting between layers [41]. Bias-sputtered SiO_2 has been an adequate interlevel dielectric for up to three levels of wiring at minimum pitch of about $2.6 \mu\text{m}$ [66] and has been used to demonstrate fabrication of stacked junction arrays [68], a gate array of 10 584 gates [78], and an 8-b 20-GHz microprocessor chip containing 63 000 junctions [17]. However, as increased circuit densities require submicrometer features and additional wiring layers, alternative global planarization techniques will be needed to achieve higher manufacturing yield.

D. Planarization

In addition to bias-sputtered SiO_2 discussed in the previous section, many other planarization techniques are available for superconductor integrated circuit fabrication. Suitable planarization techniques include liftoff and etchback planarization [79], [80], photoresist etchback planarization [81], spin-on planarizing polymer [82], mechanical polishing planarization (MPP) [83], chemical mechanical planarization (CMP) [84]–[86], and anodization [87]. All of these planarization techniques have been used with some degree of success, but none has demonstrated sufficient manufacturing yields with the possible exception of the anodization process.

CMP has gained wide acceptance by the semiconductor industry and is capable of planarizing both oxide and metal. However, the thicknesses of the oxide interlevel dielectric layers in superconductor integrated circuits tend to be thinner than in semiconductor integrated circuits. For the CMP process, this requires tight control over global planarity and precise etch stop [86]. In contrast to CMP, which uses an alkaline slurry, the MPP process uses a neutral slurry to slow down the polishing rate and a thin niobium film as an etch stop [83]. The MPP process improves control of oxide thicknesses and may become the planarization process of choice for VLSI and ULSI superconductor integrated circuit fabrication. Metal CMP or MPP could be used to form superconducting niobium “plugs” or vertical “pillars” for high-density contacts between wiring layers, but this process has yet to be demonstrated [83].

An attractive alternative to oxide interlevel dielectric is the photosensitive polyimide interlevel dielectric material described in [82]. This material is applied to the wafer by spin coating and can be partially planarizing under the appropriate spin-on conditions. The photosensitive polyimide film is exposed on a i-line stepper and developed in an alkaline solution to form contact holes. After develop, the polyimide interlevel dielectric layer only needs a 10-min 150°C bake to drive off solvents before depositing the next niobium wiring layer. This process eliminates the extra photoresist step required to define contacts in oxide interlevel dielectric. Multiple fine pitch wiring layers and submicrometer feature definition have yet to be demonstrated in this material, but photosensitive polyimide could prove to be an excellent passivation layer for superconductor integrated circuits.

In the NGST 8-kA/cm^2 process, features etched in the niobium ground plane, which are typically moats or trenches designed to prevent flux trapping, are planarized partially using a niobium anodization process that is distinct from that described by Kircher [87]. The degree of planarization achievable by this process is over 50%. It has excellent global uniformity and has demonstrated very high manufacturing yield. Fig. 11 illustrates the niobium ground plane anodization process. The niobium ground plane deposition is split into two depositions. The first niobium deposition is typically 100 nm. Next, the ground plane is patterned and etched to create ground etch features as shown in Fig. 11(a). The photoresist mask is stripped, and a second thinner niobium




Fig. 11. Ground plane planarization process. (a) Deposit first 100 nm of the niobium ground plane, mask, and etch to define circuit features. (b) Strip photoresist and deposit second 50 nm of the niobium ground plane. (c) Mask ground contacts and anodize niobium ground plane to the desired thickness (144 nm). (d) Completed process showing 52% reduced step height (112 nm) compared to ground etch step height if unfilled (235 nm).

film (typically 50 nm) is deposited, which completely covers the ground etch layer as shown in Fig. 11(b). The total niobium ground plane thickness is 150 nm, and the ground etch areas are filled with 50 nm of Nb. Next, the ground etch areas are masked, and the entire niobium ground layer including the niobium in the ground etch areas is anodized. The thinner niobium in the ground etch areas, deposited by the second deposition, is converted completely to Nb₂O₅ (typically 122 nm thick) before the desired thickness of Nb₂O₅ on the ground plane is reached (typically 144 nm). Fig. 11(c) illustrates the anodization of the ground plane and ground etch areas. The completed structure after anodization and photoresist strip is shown Fig. 11(d) and illustrates the reduction in step height from 235 nm without ground etch oxide fill to 112 nm with ground etch filled with Nb₂O₅. In this case, the degree of planarization is about 52%, and typical across-wafer step height variation is on the order of only $\pm 3\%$.

Planarization by anodization is a relatively simple process and requires only a minor change to the standard niobium ground plane anodization step. As shown in Fig. 12, this process has dramatically reduced electrical shorts, as measured by comb-to-meander test structures, between adjacent wires over ground etch steps. With the exception of a few random electrical shorts due to defects (particles), electrical

shorts in the first wiring layer were eliminated, and the reduction in step height has increased wire critical current by $\sim 74\%$ to 40 mA/ μm . Fig. 13 is a SEM picture of part of a series-biased circuit showing the high quality of the first and second wiring layers crossing over an Nb₂O₅ oxide-filled moat to connect to a circuit on an isolated ground plane.

E. Resistor Fabrication and Parameter Spreads

As the J_c increases, higher resistivity materials are required for shunted junctions to minimize circuit parasitics in RSFQ circuits. Attractive materials are sputter-deposited thin films of MoN_x [49] or NbN_x [69] because their resistivity can be adjusted over a wide range by varying the amount of nitrogen. Both materials are easily dry etched in SF₆ using existing RIE tools and recipes. In the NGST 8-kA/cm² process, a MoN_x film, adjusted to 5.0 Ω/square [18], is used for shunting junctions and biasing. The 8-kA/cm² process also includes a 0.15 Ω/square Mo/Al bilayer film [18] that is used for extremely low value shunts or for breaking a superconducting loop. At the present level of integration, both resistors have acceptable within-wafer sheet resistance spreads of 2.9% (1σ) for MoN_x and 3.2% (1σ) for Mo/Al. The spreads are almost entirely due the spatial variation in film thickness and could be reduced substantially by improving deposition geometry and sputter gun uniformity. Important resistor parameters and spreads are summarized in Table 7. Wafer-to-wafer variation of sheet resistance, as measured by standard Van der Pauw structure, is less than 6% for both resistors and indicates that these resistors have good process stability. The next-generation 20-kA/cm² process uses NbN_x resistors [69] targeted to an optimum 8.0 Ω/square . Preliminary results suggest that NbN_x resistor parameter spreads and run-to-run stability are comparable to MoN_x.

V. INTEGRATED CIRCUIT MANUFACTURABILITY

In order to produce working integrated circuits of any reasonable size, a reliable process is needed. In order to establish reliability, the foundry process capability needs to be understood. This understanding derives from application of techniques such as statistical process control (SPC), which can be used to track long-term behavior in the process. Proven in many manufacturing industries, including semiconductor manufacturing, SPC and design of experiments (DOE) improve efficiency and effectiveness in problem solving and process maintenance efforts. SPC is a powerful tool that encompasses a wide range of statistical techniques, including control charts, Pareto charts, and cause and effect diagrams [88]. The goal of SPC is to determine the inherent process variation, identify common-cause versus special-cause variation, set realistic parameter specifications, and prevent processes from going out of control while working to reduce the inherent process variability. The process variability can then be reduced with techniques such as DOE. We discuss the application of these tools to superconductor integrated circuit manufacturing.




Fig. 12. Trend chart of leakage conductance for comb-to-meander structure showing that oxide fill eliminates shorts between adjacent comb-to-meander wires over ground etch meander in PCM test structure. Wire pitch = 2.6 μm . Ground etch pitch = 4.0 μm .




Fig. 13. SEM picture of first and second wiring layers crossing ground moat filled with oxide.

Table 7
Resistor Parameters and Spreads

	Resistor	Mo/Al	MoN _x
Thickness		25 / 69 nm	95 nm
Sheet Resistance		0.15 $\Omega/\text{sq.}$	5.0 $\Omega/\text{sq.}$
Within-Wafer Spread (1σ)		3.2%	2.9%
1 st Order Gradient / cm		1.8%	2.3%
Electrical Linewidth Variation (1σ)		0.07 μm	0.03 μm

A. Process Stability

SPC techniques have been applied for in-process monitoring and for tracking important device performance on completed wafers. The purpose of in-process tracking is to catch errors as early as possible, to provide a source of information to identify where improvements are needed, and for troubleshooting when problems arise. In-process monitoring of resistor sheet resistance, film thickness, and feature line-size for critical layers has been routinely performed.

To track parametric performance, many groups use standard parametric control monitor (PCM) test chips, which are fabricated and tested on every wafer [20], [89]–[91]. Computer-aided testing provides the quantity of data required to establish a reliable process database and provides rapid and accurate feedback to the fabrication line. At NGST, more than 100 device parameters were measured including junction gap voltage, junction critical currents, junction leakage current, contact and wire critical currents, capacitor breakdown, niobium/resistor contact resistance, and inductance per square. Each parameter was tracked in a separate trend chart, such as is shown in Fig. 3 for J_c . These data were used to establish the design rules, which provided the circuit designers the range of parameters to be expected on any given chip. This is one important criterion for an integrated circuit foundry process [28], [66], [90].

Extensive data on the parameter spreads for important circuit elements such as junctions, resistors, and inductors, indicate that parameter variations for the present-day processes are $\sim 2\text{--}4\%$ (1σ) for local variations (on-chip), $\sim 5\%$ (1σ) for global variation (across-wafer), and $\sim 10\%$ (1σ) for run-to-run reproducibility [67]. Local spreads are the most important in determining circuit size; global variation will limit yield (i.e., number of good chips), as will run-to-run targeting. Present-day spreads are consistent with gate densities of ~ 100 k gates/ cm^2 , excluding other yield-limiting effects such as defect density [16].

DOE or experimental design has become an accepted framework for process development, optimization, and reduction of process variance in many manufacturing areas [92], [93]. It is an efficient tool for determining which factors affect the outcome of a process. From a cost and time point of view, it can be beneficial to vary more than one factor at a time to determine how factors affect the process. DOE has been used at NGST and by others in the superconductor electronics community for over a decade in process development activities [20], [41], [94]. These studies have enabled rapid progress in developing stable, robust processes.

B. Yield

Yield measurements are an important part of any manufacturing process [95]. Several parts, including wafer fabrication yield, parametric test yield, functional test yield, and packaging yield factor in overall product yield. Of these, wafer fabrication yield is the easiest to measure. Parametric test yield is being addressed through the use of standard test vehicles as described previously. Quantitative estimates of functional test yield may be difficult because it requires enough resources to test a large quantity of circuits of a given type to establish reliable yield statistics. Ideally, measured circuit yield can be correlated with parametric test yield and used to project yield of more complex circuits [96]–[98]. Testing at speed is also important because margins may be strong functions of frequency over some range. Packaging yield is important in the production of products or systems and will become more important as the superconductor electronics community fields more products based on active circuits.

Although parameter variations are known, little work has been done in superconductor electronics in the area of yield assessment. Since it is difficult to determine circuit yield from measurements of discrete device components, one would ideally develop a product-oriented yield vehicle that will allow a better understanding of the process yield, to correlate its yield with PCM measurements of spreads, etc., and to predict yield on future products. Several yield vehicles such as RAM [99] and shift registers [91], [98] have been developed to provide this type of information.

In the absence of more extensive yield data, one can nevertheless make projections about the effects of defect density and its effect on die yield by leveraging the yield models developed for the semiconductor industry. These models are directly applicable to superconductor electronics because the fabrication processes are similar. The relationship between yield and defect density is based on the model $Y = (1 + \beta AD)^{-1/\beta}$, where A is the defect sensitive area, D is defect density, and β describes how the defects tend to cluster on the wafer [95]. For a given defect density, smaller area chips will have a higher yield than larger area chips. Estimates for present defect densities are 1 or 2 defects/cm² and $\beta \sim 0.5$ (defects clustered toward the edges of the wafer). Based on this model, it is clear that defect density is an important consideration for high yield. NGST concluded, after a major effort to reduce gross visual defects, that most defects were induced by process tools and not by the fabrication facilities that were already operating at class 10 or better. Further reductions in defect density would require cleaner process tools.

VI. PROCESS BENCHMARKING: STATIC DIVIDER PERFORMANCE

The maximum operating speed of the toggle flip-flop (TFF) has become the standard measure or benchmark used to compare superconductor integrated circuit fabrication processes [18], [100], [101]. The TFF also has been used to compare the performance of semiconductor processes




Fig. 14. Maximum reported TFF divider speed f_{MAX} versus J_c for trilayers from HYPRES, NGST, and SUNY. The numbers adjacent to the NGST points indicate the optimum $I_c R_N$ product of the shunted junctions used in the TFF. Also, shown is the projected divider speed of ~ 450 GHz for the next-generation 20-kA/cm² process.

[102], but it is important to distinguish between true “static” divide-by-two operation from narrow band operation, which often results in inflated claims of switching speed. True static divide-by-two operation means that a well-designed TFF operates correctly from near dc to its maximum reported frequency without adjusting its bias point. The near-dc frequency response is necessary if the gate is to function with arbitrary data which may have long strings of logical zeros or ones. Superconductor integrated circuit fabrication process benchmarks or maximum TFF speeds are based on the more stringent measurements of static divider performance.

The NGST standard benchmark circuit is a 12-stage static divider that consists of an on-chip voltage-controlled oscillator (VCO) (a dc SQUID with damped junctions) and a 12-b TFF counter chain. Each stage of the counter chain uses of a symmetric four-junction TFF with symmetric current bias and a separate magnetic flux bias, which is described elsewhere [103]. The last two bits of the counter chain have self-resetting junction outputs that can be counted by a room temperature electronic frequency counter.

The circuit parameters are chosen to optimize operating margin and yield at high speed. The design of the NGST static divider uses junctions that are slightly underdamped.¹

The maximum divider speeds achieved are just above 200 and 300 GHz for the 4-kA/cm² and 8-kA/cm² processes, respectively. These results along with results from HYPRES and SUNY for other J_c 's are shown in Fig. 14. The maximum divider speed (f_{MAX}) scales approximately as $[J_c \text{ (kA/cm}^2\text{)}]^{1/2} \times 100$ GHz to about 50 kA/cm², above which the speed saturates at a frequency corresponding to the gap frequency for Nb. There is good agreement among the different niobium trilayer processes shown in Fig. 14. Details of the divider speed measurements and characterization can be found in [18], [59].

¹For the 8-kA/cm² process, $I_c R_N = 1.05$ mV for a Stuart–McCumber parameter (β_C) = 2.5 and for 4-kA/cm², $I_c R_N = 0.7$ mV, and $\beta_C = 2.0$.




Fig. 15. Typical current-voltage characteristics of a 20-kA/cm², 0.90-μm (drawn diameter) junction. Low subgap leakage (lower curve) is revealed when I_c is suppressed by magnetic field. $I_c \sim 60 \mu\text{A}$, $V_m \sim 13 \text{ mV}$, and $V_{\text{gap}} \sim 2.7 \text{ mV}$. Electrical sizing = $0.61 \mu\text{m}$ (effective diameter).




Fig. 16. Plot and linear fit of the square root- I_c versus drawn (e-beam written) junction diameter for a series of 20-kA/cm² junctions ranging in diameter from $0.7 \mu\text{m}$ to $1.4 \mu\text{m}$. Slope of the linear fit indicates J_c , and the x-intercept determines process bias or CD loss from the drawn feature.

VII. ADVANCED JUNCTION DEVELOPMENT

A high- J_c , submicrometer junction process was developed at NGST in collaboration with Jet Propulsion Laboratory (JPL) using their high-definition e-beam lithography tool [104]. The typical I - V characteristics of a $0.9\text{-}\mu\text{m}$ -diameter 20-kA/cm² junction are shown in Fig. 15. The junction has an effective electrical size of $0.61 \mu\text{m}$ and $V_m \sim 13\text{--}14 \text{ mV}$. The shrink or CD loss is attributed to process bias, which is equal to the x-intercept (but opposite sign) from the linear fit of the square root- I_c versus drawn (e-beam written) junction diameters. The plot of actual junction data and the fit to the data are shown in Fig. 16 for the range of sizes from 0.7 to $1.4 \mu\text{m}$. The fit is excellent as indicated by a nearly ideal correlation coefficient of 0.9998 (1.0 for perfect fit). Similar results from several measurements indicate that within-wafer uncertainty in junction sizing (variation in x-intercept in Fig. 16) is on the order of $\pm 0.01 \mu\text{m}$. This result is equal to the best results obtained in the 8-kA/cm² process for larger $1.25\text{-}\mu\text{m}$ junctions that were defined using

a g-line 1X stepper and the lithography process discussed in Section IV.

The junctions at 20 kA/cm^2 are well behaved and do not show unusual transport properties, consistent with the view discussed in [11]. These junctions have acceptable I - V characteristics for RSFQ logic and have demonstrated consistent electrical sizing. I_c spreads from the first $0.8\text{-}\mu\text{m}$ 100-junction arrays fabricated in the 20-kA/cm^2 process were on the order of 2.3% (1σ), which is a very encouraging result. However, more junction development remains to demonstrate lower on-chip I_c spreads ($\sigma < 2\%$), J_c targeting, and manufacturing yields, since at 20 kA/cm^2 , the oxidation process is well into the high- J_c , low exposure regime where J_c is more sensitive to variations in the oxidation conditions [58]. Based on these junction results, there does not appear to be any fundamental physical roadblock to impede deployment of the next-generation process.

VIII. ADVANCED PROCESS DEVELOPMENT

The next-generation niobium-based process should target 20 kA/cm^2 minimum size junctions of $0.8 \mu\text{m}$ and six metal layers. RSFQ logic circuits fabricated in this process are expected to achieve on-chip clock rates of $80\text{--}100 \text{ GHz}$. Present scaling models predict that gate densities can be greater than $50\,000 \text{ gates/cm}^2$ [105]. The projected maximum TFF divider speed is 450 GHz as indicated in Fig. 14. A 20-kA/cm^2 process should build upon the foundation and success of the NGST 8-kA/cm^2 process and include two additional wiring layers and advanced process features including ground plane planarization by anodization and light junction anodization. The additional wiring layers will need to be planarized. As the process integration matures, all lithography steps, not only the critical junction and contact definition steps, should transition to at least an i-line stepper capable of $0.35\text{-}\mu\text{m}$ resolution to reduce all feature sizes.

While the next-generation RSFQ circuits are expected to operate at higher speeds, perhaps the most significant aspect in moving to the 20-kA/cm^2 process is the reduction in chip size and increase in gate density. For example, the 8-b 20-GHz microprocessor chip (FLUX-1r1), designed in the NGST 4-kA/cm^2 process, occupies a chip area of 1 cm^2 (Fig. 17). If the chip is redesigned for a 20-kA/cm^2 process utilizing the additional wiring layers, planarization, and feature size reductions, the circuit could potentially occupy only about one-ninth of the area (see inset in Fig. 17) compared to the 4-kA/cm^2 design. Smaller chips will increase yield and will reduce time of flight latency to improve computing efficiency.

Niobium-based superconductor chip fabrication process has been advancing at the rate of about one generation every two years since 1998, as shown by the NGST roadmap for niobium technology in Table 8. As the roadmap indicates, evolutionary, not revolutionary, improvements are needed to realize the full potential of niobium technology. Comparison with the Semiconductor Industry Association (SIA) roadmap suggests that 1995 semiconductor fabrication technology is comparable to that needed to produce superconductor chips




Fig. 17. Photograph of the FLUX-1r1 microprocessor chip fabricated in NGST's 4-kA/cm² process. The chip is 1 cm on a side, contains ~63 000 junctions, and dissipates 9 mW. The inset represents the potential reduction in chip size if the FLUX chip is redesigned in a 20-kA/cm² process.

Table 8
Nb Technology Roadmap

	Nb Technology					SIA CMOS
Year	1998	2000	2002	2004	Future	1995
Minimum feature (μm)	1.5	1.0	1.0	0.6	0.35	0.35
Minimum junction (μm)	2.5	1.75	1.25	0.8	0.5	N/A
J_c (kA/cm ²)	2	4	8	20	50	N/A
Number of masks	12	12	14	18	22	~ 20
Resistors	2	2	2	2	2	N/A
Metal line pitch (μm)	4.0	3.0	2.6	2.0	1.0	0.9
Metal interconnect layers	4	4	4	6	8	3
Planarization	no	no	partial	yes	yes	yes
Max. defects per cm ²	2	2	1	1	.1	.1
I/O count	128	128	128	640	1200	900
F_{max} (GHz), divider freq.	150	210	300	450	750	1.5
Processor Clock (GHz)	30	40	60	80	150	0.3
Gate Density (gates/cm ²)	5K	8K	16K	64K	>128K	300K

for petaFLOPS scale computing. Semiconductor technology has scaled by ~ 0.7 in feature size and ~ 2.5 in gate density for each generation, each of which has taken about three years and a substantial capital investment. Beyond the present 8- to 10-kA/cm² processes, continued evolutionary improvements in niobium technology will require only relatively modest investment in new process tools and facilities.

IX. CONCLUSION

Because of its extreme speed advantage *and* ultralow power dissipation, superconductor electronics based on

RSFQ logic will eventually find widespread use in high-performance computing and in applications that need the fastest digital technology such as analog-to-digital converters and digital signal processing. To be a viable commercial digital technology, the superconductor electronics community must demonstrate *systems* that operate at higher speeds with lower power than any other technology. Process development efforts are pushing the boundaries of niobium technology in three key areas: increasing speed, increasing complexity, and increasing yields. Rapid progress has been achieved in the past few years, driven by the high-performance computing initiative. In addition to advancements in fabrication, the superconductor electronics community needs to focus attention on standardization of test and design tools. Finally, building complete systems or subsystems, which have taken into account I/O, magnetic shielding, cooling, and other packaging issues will be required to expand the customer base beyond a few specialty applications.

ACKNOWLEDGMENT

The authors would like to thank the NGST foundry team for their many years of fabrication support, particularly in establishing and maintaining stable fabrication processes, and the NGST test and packaging teams for their many years of timely support with parametric data, trend charting, and wafer preparation. The authors would also like to thank A. Kleinsasser and the e-beam lithography team at JPL for a very successful collaboration, and F. D. Bedard of DoD and D. Van Vechten of ONR for their continuing support of advanced process development, and John Spargo, Andy Smith, and Paul Bunyk of NGST and Professor Ted Van Duzer of University of California, Berkeley for their valuable discussions. We especially want to acknowledge the many individuals who contributed to Table 1: K. K. Berggren, Massachusetts Institute of Technology; F.-Im. Buchholz, Physikalisch-Technische Bundesanstalt; W. Chen, State University of New York, Stony Brook; P. Dresselhaus, National Institute of Standards and Technology (NIST); M. Hidaka, SRL; G. Hilton, NIST; Prof. W. Jutzi, Universität Karlsruhe; X. Meng, University of California, Berkeley; H.-G. Meyer, IPHT-Jena; A. Shoji, AIST; S. K. Tolpygo, HYPRES; and J.-C. Villégier, CAE-Grenoble.

REFERENCES

- [1] K. K. Likharev and V. K. Semenov, "RSFQ logic/memory family: A new Josephson-junction technology for sub-terahertz-clock-frequency digital systems," *IEEE Trans. Appl. Superconduct.*, vol. 1, pp. 3–28, Mar. 1991.
- [2] W. Chen, A. V. Rylyakov, V. Patel, J. E. Lukens, and K. K. Likharev, "Rapid single flux quantum T-flip flop operating up to 770 GHz," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3212–3215, June 1999.
- [3] K. K. Likharev, "RSFQ: The fastest digital technology," in *Proc. 5th Eur. Workshop Low Temperature Electronics 2002*, p. 155.
- [4] T. Sterling, "How to build a hypercomputer," *Sci. Amer.*, pp. 38–45, July 2001.
- [5] T. Sterling, G. Gao, K. K. Likharev, P. M. Kogge, and M. J. MacDonald, "Steps to petaflops computing: A hybrid technology multi-threaded architecture," in *Proc. IEEE Aerospace Conf. 1997*, vol. 2, pp. 41–75.

- [6] K. B. Theobald, G. R. Gao, and T. L. Sterling, "Superconducting processors for HTMT: Issues and challenges," in *Proc. 7th Symp. Frontiers of Massively Parallel Computation 1999*, pp. 260–267.
- [7] K. K. Likharev, "Superconductor devices for ultrafast computing," in *Applications of Superconductivity*, H. Weinstock, Ed. Dordrecht, The Netherlands: Kluwer, 1999.
- [8] International Technology Roadmap for Semiconductors: Emerging research devices (2003). [Online]. Available: <http://public.itrs.net/Files/2003ITRS/ERD2003.pdf>
- [9] M. Tanaka, F. Matsuzaki, T. Kondo, N. Nakajima, Y. Yamanashi, A. Fujimaki, H. Hayakawa, N. Yoshikawa, H. Terai, and S. Yorozu, "A single-flux-quantum logic prototype microprocessor," in *Dig. Tech. Papers 2003 IEEE Int. Solid-State Circuits Conf.*, pp. 1–3.
- [10] J. Spargo, Q. Herr, and L. Abelson, "High speed digital devices, circuits, and applications in the U.S.," in *Extended Abstracts 8th Int. Superconductivity Electronics Conf.*, 2001, pp. 101–102.
- [11] P. Bunyk, K. Likharev, and D. Zinoviev, "RSFQ technology: Physics and devices," *Int. J. High Speed Electron. Syst.*, vol. 11, pp. 257–305, 2001.
- [12] Y. Makhlin, G. Schön, and A. Shnirman, "Quantum-state engineering with Josephson-junction devices," *Rev. Mod. Phys.*, vol. 73, pp. 357–400, Apr. 2001.
- [13] V. K. Semenov and D. V. Averin, "SFQ control circuits for Josephson junction qubits," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 960–965, June 2003.
- [14] M. F. Bocko, A. M. Herr, and M. J. Feldman, "Prospects for quantum coherent computation using superconducting electronics," *IEEE Trans. Appl. Superconduct.*, vol. 7, pp. 3638–3641, June 1997.
- [15] A. W. Kleinsasser, "High performance Nb Josephson devices for petaflops computing," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 1043–1049, Mar. 2001.
- [16] L. A. Abelson, Q. P. Herr, G. L. Kerber, M. Leung, and T. S. Tighe, "Manufacturability of superconductor electronics for a petaflops-scale computer," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3202–3207, June 1999.
- [17] A. Silver, A. Kleinsasser, G. Kerber, Q. Herr, M. Dorojevets, P. Bunyk, and L. Abelson, "Development of superconductor electronics technology for high end computing," *Supercond. Sci. Technol.*, vol. 16, pp. 1368–1374, Oct. 2003.
- [18] G. L. Kerber, L. A. Abelson, K. Edwards, R. Hu, M. W. Johnson, M. L. Leung, and J. Luine, "Fabrication of high current density Nb integrated circuits using a self-aligned junction anodization process," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 82–86, June 2003.
- [19] Q. P. Herr, M. S. Wire, and A. D. Smith, "High speed data link between digital superconductor chips," *Appl. Phys. Lett.*, vol. 80, pp. 3210–3212, Apr. 2002.
- [20] J. E. Savageau, C. J. Burroughs, P. A. A. Booi, M. W. Cromar, S. P. Benz, and J. A. Koch, "Superconducting integrated circuit fabrication with low temperature ECR-based PECVD SiO₂ dielectric films," *IEEE Trans. Appl. Superconduct.*, vol. 5, pp. 2303–2309, June 1995.
- [21] K. K. Berggren, E. M. Macedo, D. A. Feld, and J. P. Sage, "Low T_c superconductive circuits fabricated on 150-mm-diameter wafers using a doubly planarized Nb/AlO_x/Nb process," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3271–3274, June 1999.
- [22] R. Dolata, M. I. Khabipov, F.-Im. Buchholz, W. Kessel, and J. Niemeyer, "Nb/Al₂O₃-Al/Nb process development for the fabrication of fast-switching circuits in RSFQ logic," in *Proc. 2nd Eur. Conf. Applied Superconductivity (EUCAS'95)*, vol. 2, pp. 1709–1712.
- [23] D. Balashov, M. I. Khabipov, F.-Im. Buchholz, W. Kessel, and J. Niemeyer, "SINIS fabrication process for realizing integrated circuits in RSFQ impulse logic," *Supercond. Sci. Technol.*, vol. 12, pp. 864–867, 1999.
- [24] M. I. Khabipov, D. Balashov, F.-Im. Buchholz, W. Kessel, and J. Niemeyer, "RSFQ circuitry realized in a SINIS technology process," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 4682–4687, Dec. 1999.
- [25] X. Meng, A. Wong, and T. Van Duzer, "Micron and submicron Nb/Al-AlO_x-Al/Nb tunnel junctions with high critical current densities," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 365–368, Mar. 2001.
- [26] M. Maezawa, M. Suzuki, and A. Shoji, "Design and operation of a rapid single flux quantum demultiplexer," *Supercond. Sci. Technol.*, vol. 15, pp. 1744–1746, 2002.
- [27] H. Yamamori, M. Itoh, H. Sasaki, S. P. Benz, and P. D. Dresselhaus, "All-NbN digital-to-analog converters for a programmable voltage standard," *Supercond. Sci. Technol.*, vol. 14, p. 1048, Dec. 2001.
- [28] HYPRES design rules. HYPRES, Elmsford, NY. [Online]. Available: <http://www.hypres.com>
- [29] S. Topylgo, private communication, Nov. 2003.
- [30] J. C. Villegier, M. Salez, F. Miletto-Granozio, and A. Karpov, "Processing and characterization of high Jc NbN superconducting tunnel junctions for THz analog circuits and RSFQ," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3216–3219, June 1999.
- [31] G. L. Kerber, L. A. Abelson, R. N. Elmadjian, and E. G. Ladizinsky, "Characteristics of junctions and resistors fabrication using an all-NbN superconductor integrated circuit foundry process," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3267–3270, June 1999.
- [32] W. Chen, V. Patel, S. K. Tolpygo, D. Yohannes, S. Pottorf, and J. E. Lukens, "Development toward high-speed integrated circuits and SQUID qubits with Nb/AlO_x/Nb Josephson junctions," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 103–106, June 2003.
- [33] S. Tahara, S. Yorozu, Y. Kameda, Y. Hashimoto, H. Numata, T. Satoh, W. Hattori, and M. Hidaka, "Superconducting digital electronics," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 463–468, Mar. 2001.
- [34] L. Fritzsche, H. Elsner, M. Schubert, and H.-G. Meyer, "SNS and SIS Josephson junctions with dimensions down to the sub-micron region prepared by a unified technology," *Supercond. Sci. Technol.*, vol. 12, pp. 880–882, 1999.
- [35] S. P. Benz, "Superconductor-normal-superconductor junctions for programmable voltage standards," *Appl. Phys. Lett.*, vol. 67, pp. 2714–2716, Oct. 1995.
- [36] R. Lochschmied, R. Herwig, M. Neuhaus, and W. Jutzi, "A low power 12 bit flux shuttle shift register with Nb technology," *IEEE Trans. Appl. Superconduct.*, vol. 7, pp. 2983–2986, June 1997.
- [37] H. Terai and Z. Wang, "9 K operation of RSFQ logic cells fabricated by NbN integrated circuit technology," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 525–528, Mar. 2001.
- [38] O. Kieler, R. Behr, F. Müller, H. Schulze, J. Kohlmann, and J. Niemeyer, "Improved 1 V programmable Josephson voltage standard using SINIS junctions," *Physica C*, vol. 372–376, pp. 309–311, 2002.
- [39] J. K. Freericks, B. K. Nikolic, and P. Miller, "Superconductor-correlated metal-superconductor Josephson junctions: An optimized class for high speed digital electronics," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 1089–1092, June 2003.
- [40] I. P. Nevirkovets, S. E. Shafranjuk, J. B. Ketterson, and E. M. Rudenko, "Properties of high-jc SINIS junctions," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 1085–1088, June 2003.
- [41] L. Abelson, "Superconductive electronics process technologies," in *Extended Abstracts 6th Int. Superconductivity Electronics Conf.*, June 1997, pp. 1–4.
- [42] X. Meng, A. Bhat, and T. Van Duzer, "Very small critical current spreads in Nb/AlO_x/Nb integrated circuits using low temperature and low stress ECR PECVD silicon oxide films," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3208–3211, June 1999.
- [43] A. Shoji, M. Aoyagi, S. Kosaka, and F. Shinoki, "Temperature-dependent properties of niobium nitride Josephson tunnel junctions," *IEEE Trans. Magn.*, vol. 23, pp. 1464–1471, Mar. 1987.
- [44] M. Radparvar, L. S. Yu-Jahnes, and R. T. Hunt, "All niobium nitride Josephson tunnel junctions with thermally oxidized magnesium barrier," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2050–2053, Mar. 1993.
- [45] S. Kubo, M. Asahi, M. Hikita, and M. Igarashi, "Magnetic penetration depths in superconducting NbN films prepared by reactive dc magnetron sputtering," *Appl. Phys. Lett.*, vol. 44, pp. 258–260, Jan. 15, 1984.
- [46] S. L. Thomasson, A. W. Moopenn, R. Elmadjian, J. M. Murduck, J. W. Spargo, L. A. Abelson, and H. W. Chan, "All refractory NbN integrated circuit process," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2058–2061, Mar. 1993.

- [47] B. Bumble, H. G. LeDuc, J. A. Stern, and K. G. Megerian, "Fabrication of Nb/Al-Nx/NbTiN junctions for SIS mixer applications," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 76–79, Mar. 2001.
- [48] M. Gurvitch, W. A. Washington, and H. A. Huggins, "High quality refractory Josephson tunnel junctions utilizing thin aluminum layers," *Appl. Phys. Lett.*, vol. 42, pp. 472–474, Mar. 1983.
- [49] Y. Tarutani, M. Hirano, and U. Kawabe, "Niobium-based integrated circuit technologies," in *Proc. IEEE*, vol. 77, Aug. 1989, pp. 1164–1176.
- [50] T. Imamura, T. Shiota, and T. S. Hasuo, "Fabrication of high quality Nb/AlO_x-Al/Nb Josephson junctions. I. Sputtered Nb films for junction electrodes," *IEEE Trans. Appl. Superconduct.*, vol. 2, pp. 1–14, Mar. 1992.
- [51] T. Imamura and T. S. Hasuo, "Fabrication of high quality Nb/AlO_x-Al/Nb Josephson junctions. II. Deposition of thin Al layers on Nb films," *IEEE Trans. Appl. Superconduct.*, vol. 2, pp. 84–94, June 1992.
- [52] —, "Fabrication of high quality Nb/AlO_x-Al/Nb Josephson junctions. III. Annealing stability of AlO_x tunneling barriers," *IEEE Trans. Appl. Superconduct.*, vol. 2, pp. 222–227, Dec. 1992.
- [53] T. Imamura and S. Hasuo, "Effects of intrinsic stress on submicrometer Nb/AlO_x/Nb Josephson junctions," *IEEE Trans. Magn.*, vol. 25, pp. 1119–1122, Mar. 1989.
- [54] K. Tsukada, J. Kawai, G. Uehara, and H. Kado, "Relationship of Nb surface morphology and Al coverage to the intrinsic stress of magnetron-sputtered Nb films," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2944–2946, Mar. 1993.
- [55] P. A. A. Booi, C. A. Livingston, and S. P. Benz, "Intrinsic stress in DC sputtered niobium," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 3029–3031, June 1993.
- [56] R. S. Amos, P. E. Breyer, H. H. Huang, and A. W. Lichtenberger, "Stress and source conditions of DC magnetron sputtered Nb films," *IEEE Trans. Appl. Superconduct.*, vol. 5, pp. 2326–2329, June 1995.
- [57] K. Kuroda and M. Yuda, "Niobium-stress influence on Nb/Al-oxide/Nb Josephson junctions," *J. Appl. Phys.*, vol. 63, pp. 2352–2357, Apr. 1, 1988.
- [58] A. W. Kleinsasser, R. E. Miller, and W. H. Mallison, "Dependence of critical current density on oxygen exposure in Nb-AlO_x-Nb tunnel junctions," *IEEE Trans. Appl. Superconduct.*, vol. 5, pp. 26–30, Mar. 1995.
- [59] G. L. Kerber, L. A. Abelson, M. L. Leung, Q. P. Herr, and M. W. Johnson, "A high density 4 kA/cm² Nb integrated circuit process," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 1061–1065, Mar. 2001.
- [60] M. Aoyagi, M. Maezawa, H. Nakagawa, and I. Kurosawa, "Fabrication of submicron Nb/AlO_x/Nb Josephson junctions using ECR plasma etching technique," *IEEE Trans. Appl. Superconduct.*, vol. 7, pp. 2644–2648, June 1997.
- [61] H. Kroger, L. N. Smith, and D. W. Jillie, "Selective niobium anodization process for fabricating Josephson junctions," *Appl. Phys. Lett.*, vol. 39, pp. 280–282, Aug. 1, 1981.
- [62] S. Morohashi, S. Hasuo, and T. Yamaoka, "Self-aligned contact process for Nb/Al-AlO_x/Nb Josephson junctions," *Appl. Phys. Lett.*, vol. 48–3212, pp. 254–256, Jan. 20, 1986.
- [63] L. P. S. Lee, E. R. Arambula, G. Hanaya, C. Dang, R. Sandell, and H. Chan, "RHEA (resist-hardened etch and anodization) process for fine-geometry Josephson junction fabrication," *IEEE Trans. Magn.*, vol. 27, pp. 3133–3136, Mar. 1991.
- [64] X. Meng, L. Zheng, A. Wong, and T. Van Duzer, "Micron and submicron Nb/Al-AlO_x/Nb tunnel junctions with high critical current densities," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 365–368, Mar. 2001.
- [65] AZ-300T photoresist stripper and AZ 3318D photoresist. Clariant Corp., Somerville, NJ. [Online]. Available: <http://azwebsale@clariant.com>
- [66] "8-kA/cm² topological design rules, JJ110E, Rev. 2.2," NGST, Redondo Beach, CA.
- [67] L. Abelson, K. Daly, N. Martinez, and A. D. Smith, "LTS Josephson critical current densities for LSI applications," *IEEE Trans. Appl. Superconduct.*, vol. 5, pp. 2727–2730, June 1995.
- [68] S. Hasuo, "High-speed Josephson integrated circuit technology," *IEEE Trans. Magn.*, vol. 25, pp. 740–749, Mar. 1989.
- [69] G. L. Kerber, L. A. Abelson, R. N. Elmadjian, G. Hanaya, and E. G. Ladizinsky, "An improved NbN integrated circuit process featuring thick NbN ground plane and lower parasitic circuit inductances," *IEEE Trans. Appl. Superconduct.*, vol. 7, pp. 2638–2643, June 1997.
- [70] H. Numata and S. Tahara, "Fabrication technology for Nb integrated circuits," *IEICE Trans. Electron.*, vol. E84-C, pp. 2–8, Jan. 2001.
- [71] J. S. Logan, F. S. Maddocks, and P. D. Davidse, "Metal edge coverage and control of charge accumulation in rf sputtered insulators," *IBM J. Res. Develop.*, vol. 14, pp. 182–191, Mar. 1970.
- [72] J. P. Ducommun, M. Cantagrel, and M. Moulin, "Evolution of well-defined surface contour submitted to ion bombardment: Computer simulation and experimental investigation," *J. Mater. Sci.*, vol. 10, pp. 52–62, 1975.
- [73] C. Y. Ting, V. J. Vivaida, and H. G. Schaefer, "Study of planarized sputter-deposited SiO₂," *J. Vac. Technol.*, vol. 15, pp. 1105–1112, May–June 1978.
- [74] C. H. Ting and A. R. Neureuther, "Applications of profile simulation for thin film deposition and etching processes," *Solid State Technol.*, vol. 25, pp. 115–123, Feb. 1982.
- [75] A. Mumtaz, I. Drapkin, and R. Mathew, "Radio frequency magnetron sputtering of radio frequency biased quartz on a scanning pallet," *J. Vac. Sci. Technol. A*, vol. 2, pp. 237–240, Apr.–June 1984.
- [76] J. S. Logan, "Control of RF sputtered film properties through substrate tuning," *IBM J. Res. Develop.*, vol. 14, pp. 172–175, Mar. 1970.
- [77] L. I. Maissel, R. E. Jones, and C. L. Standley, "Re-emission of sputtered SiO₂ during growth and its relation to film quality," *IBM J. Res. Develop.*, vol. 14, pp. 176–181, Mar. 1970.
- [78] S. Hasuo and T. Imamura, "Digital logic circuits," *Proc. IEEE*, vol. 77, pp. 1177–1193, Aug. 1989.
- [79] S. Nagasawa, H. Tsuge, and Y. Wada, "Planarization technology for Josephson integrated circuits," *IEEE Electron Device Lett.*, vol. 9, pp. 414–416, Aug. 1988.
- [80] I. Ishida, S. Tahara, M. Hidaka, S. Nagasawa, S. Tuschida, and Y. Wada, "A fabrication process for a 580 ps 4 kbit Josephson nondestructive read-out RAM," *IEEE Trans. Magn.*, vol. 27, pp. 3113–3116, Mar. 1991.
- [81] A. T. Barfknecht, R. C. Ruby, H. L. Ko, and G. S. Lee, "Josephson junction integrated circuit process with planarized PECVD SiO₂ dielectric," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2201–2203, Mar. 1993.
- [82] K. Kikuchi, S. Segawa, E.-S. Jung, H. Nakagawa, K. Tokoro, H. Itatani, and M. Aoyagi, "New fabrication process for Josephson tunnel junctions using photosensitive polyimide insulation layer for superconducting integrated circuits," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 119–122, June 2003.
- [83] H. Numata, S. Nagasawa, M. Tanaka, and S. Tahara, "Fabrication technology for high-density Josephson integrated circuits using mechanical polishing planarization," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 3198–3201, June 1999.
- [84] M. B. Ketchen, D. Pearson, A. W. Kleinsasser, C.-K. Hu, M. Smyth, J. Logan, K. Stawiasz, E. Baran, M. Jaso, T. Ross, K. Petrillo, M. Manny, S. Basavaiah, S. Brodsky, S. B. Kaplan, and W. J. Gallagher, "Sub-μm, planarized, Nb-AlO_x-Nb Josephson process for 125 mm wafers developed in partnership with Si technology," *Appl. Phys. Lett.*, vol. 59, pp. 2609–2611, Nov. 1991.
- [85] Z. Bao, M. Bhushan, S. Han, and J. E. Lukens, "Fabrication of high quality, deep-submicron Nb/AlO_x/Nb Josephson junctions using chemical mechanical polishing," *IEEE Trans. Appl. Superconduct.*, vol. 5, pp. 2731–2734, June 1995.
- [86] J. M. Murduck, A. Kirschenbaum, A. Mayer, V. Morales, and C. Lavoie, "High-performance Nb integrated circuit process fabrication," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 87–90, June 2003.
- [87] C. J. Kircher and H. H. Zappe, "Process for fabricating above and below ground plane wiring on one side of a supporting substrate and the resulting circuit configuration," U.S. Patent 4 075 756, Feb. 1978.
- [88] D. C. Montgomery, *Introduction to Statistical Quality Control*, 2nd ed. New York: Wiley, 1991.
- [89] A. A. Joseph, S. Heuvelmans, G. J. Gerritsma, and H. G. Kerkhoff, "The detection of defects in a niobium tri-layer process," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 95–98, June 2003.

- [90] L. A. Abelson, S. L. Thomasson, J. M. Murduck, R. Elmadjian, G. Akerling, R. Kono, and H. W. Chan, "A superconductive integrated circuit foundry," *IEEE Trans. Appl. Superconduct.*, vol. 3, no. 1, pp. 2043–2049, Mar. 1993.
- [91] V. K. Semenov, Yu. A. Polyakov, and W. Chao, "Extraction of impacts of fabrication spread and thermal noise on operation of superconducting digital circuits," *IEEE Trans. Appl. Superconduct.*, vol. 9, pp. 4030–4033, June 1999.
- [92] W. Y. Fowles and C. M. Creveling, *Engineering Methods for Robust Product Design Using Taguchi Methods in Technology and Product Development*. Reading, MA: Addison-Wesley, 1995.
- [93] G. Box, W. Hunter, and J. Hunter, *Statistics for Experiments*. New York: Wiley, 1978.
- [94] J. M. Murduck, J. DiMond, C. Dang, and H. Chan, "Niobium nitride Josephson process development," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2211–2214, Mar. 1993.
- [95] A. V. Ferris-Prabhu, *Introduction to Semiconductor Device Yield Modeling*. Norwood, MA: Artech House, 1992.
- [96] D. L. Miller, J. X. Przybysz, and J. H. Kang, "Margins and yields of SFQ circuits in HTS materials," *IEEE Trans. Appl. Superconduct.*, vol. 3, pp. 2728–2731, Mar. 1993.
- [97] C. A. Hamilton and K. C. Gilbert, "Margins and yield in single flux quantum logic," *IEEE Trans. Appl. Superconduct.*, vol. 1, pp. 157–163, Dec. 1991.
- [98] D. J. Durand, L. A. Abelson, L. R. Eaton, M. Leung, A. Moopenn, J. M. Murduck, and J. Spargo, "Ten kelvin NbN integrated digital circuits," in *Extended Abstracts 4th Int. Superconductivity Electronics Conf.*, 1993, pp. 209–210.
- [99] S. Tahara, S. Nagasawa, H. Numata, Y. Hashimoto, S. Yoroazu, and H. Matsuoka, "Josephson digital LSI technology," in *Extended Abstracts 5th Int. Superconductivity Electronics Conf.*, 1995, pp. 23–25.
- [100] W. Chen, A. V. Rylyakov, V. Patel, J. E. Lukens, and K. K. Likharev, "Superconductor digital frequency divider operating up to 750 GHz," *Appl. Phys. Lett.*, vol. 73, pp. 2817–2819, Nov. 1998.
- [101] D. K. Brock, A. M. Kadin, A. F. Kirichenko, O. A. Mukhanov, S. Sawana, J. A. Vivalda, W. Chen, and J. E. Lukens, "Retargeting RSFQ cells to a submicron fabrication process," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 369–372, Mar. 2001.
- [102] R. Pulella, D. Mensa, B. Argarwal, Q. Lee, J. Guthrie, and M. J. W. Rodwell, "47 GHz static frequency divider in ultrafast transferred-substrate heterojunction bipolar transistor technology," presented at the Conf. InP and Related Materials, Tsukuba, Japan, 1998.
- [103] M. W. Johnson, B. J. Dalrymple, and D. J. Durand, "Wide bandwidth oscillator/counter A/D converters," *IEEE Trans. Appl. Superconduct.*, vol. 11, pp. 607–611, Mar. 2001.
- [104] A. Kleinsasser. E-beam lithography tool. Jet Propulsion Lab, Pasadena, CA. [Online]. Available: <http://Alan.W.Kleinsasser@jpl.nasa.gov>
- [105] P. Bunyk, "RSFQ random logic gate density scaling for the next-generation Josephson junction technology," *IEEE Trans. Appl. Superconduct.*, vol. 13, pp. 496–497, 2003.



Lynn A. Abelson received the B.S. degree in geology from University of California, Berkeley, in 1981 and the M.S. and Ph.D. degrees in geology from the California Institute of Technology, Pasadena, CA, in 1984 and 1988, respectively.

She was at the Xerox Microelectronics Center, El Segundo, CA, working in CMOS process integration. She has been Manager of Northrop Grumman Space Technology's (formerly TRW) superconductor electronics integrated circuit fabrication effort for more than 12 years.

Her research interests include integrated circuit fabrication, statistical process control, and manufacturing technology. She holds four patents in TRW's niobium and niobium nitride process technology and has over 25 publications.



George L. Kerber received the B.S. and M.S. degrees in physics from California State University, Fresno, in 1972 and 1973, respectively, and the Master of Engineering degree from the University of California, Berkeley, in 1975.

From 1975 to 1982, he worked at the Naval Ocean Systems Center, San Diego, CA, on the fabrication of various thin-film devices including surface acoustic wave filters and variable thickness microbridges for superconducting quantum interference device (SQUID) magnetometers.

From 1983 to 1987, he worked for a startup company in San Diego, where he was responsible for the development of sensors and instrumentation for production oil well logging and drilling tools. In 1987, he joined Hughes Aircraft Company, Carlsbad, CA to work on niobium nitride integrated circuit fabrication, and from 1990 to 1994, he was responsible for process integration of several CMOS and biCMOS technologies. From 1994 to 2004, he worked at Northrop Grumman Space Technology (formerly TRW), Redondo Beach, CA, and was responsible for the development and process integration of niobium nitride and niobium-based superconductor integrated circuits. He is currently a Consultant, San Diego, CA. He has 14 patents and more than 25 publications.


Appendix J

CAD

PRESENT CAD CAPABILITY

The integrated circuit design software developed by Stony Brook University and the University of Rochester was integrated into a single software suite at Northrop Grumman. The readiness of U.S.-based design methodology and CAD tools will be described in detail using the Northrop Grumman capability as an example. While built upon the academic projects, it also leverages the methodology and software that serves the commercial semiconductor ASIC world. Where possible, semiconductor standards were adhered to.

Circuit Design



RSFQ Gate Library

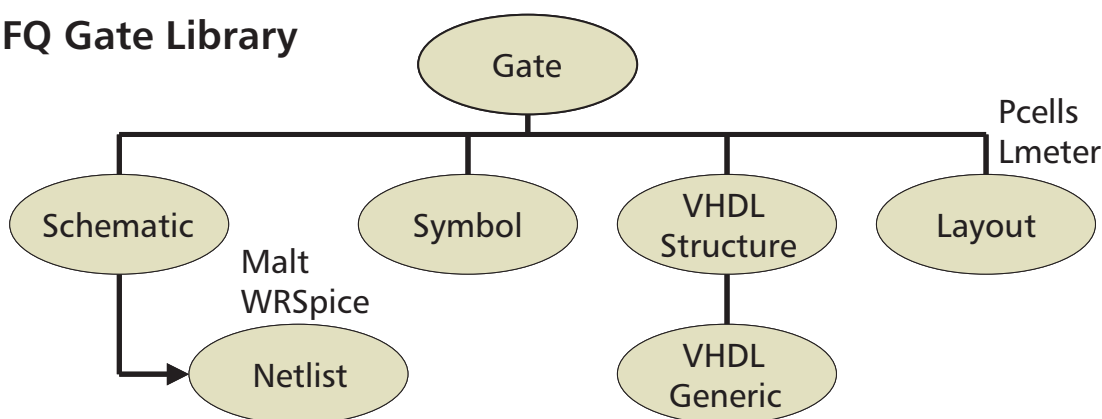


Fig. 1. Circuit design and verification, top, is distinct from logic gate library development, bottom. The chip foundry publishes the gate library to serve all users. Customers construct complex circuits from the gates, to satisfy their unique performance goals.

In the following description we will repeatedly distinguish between the gate library developer and the circuit designer. Division of labor between these two roles, illustrated in Figure 1, is among the most important steps in design methodology. The role of the circuit designer is to architect, simulate and verify complex designs. This can be accomplished without knowledge of all of the details of the device physics, foundry process, or CAD software.

As in the commercial ASIC world, device and gate libraries are published by the foundry. These include, at minimum, device models for physical-level simulation and parameterized-cell (Pcell) physical layout of devices and gates. Commercial foundries emphasize formal verification of designs submitted for fabrication. Therefore, the foundry typically provides CAD-based design verification rules for design rule checking (DRC) and layout-versus-schematic (LVS). Some commercial foundries require clean DRC and LVS verification run files to be submitted with the design prior to fabrication. The foundry may also perform additional verification checks and require design modifications to physical layout before tooling.

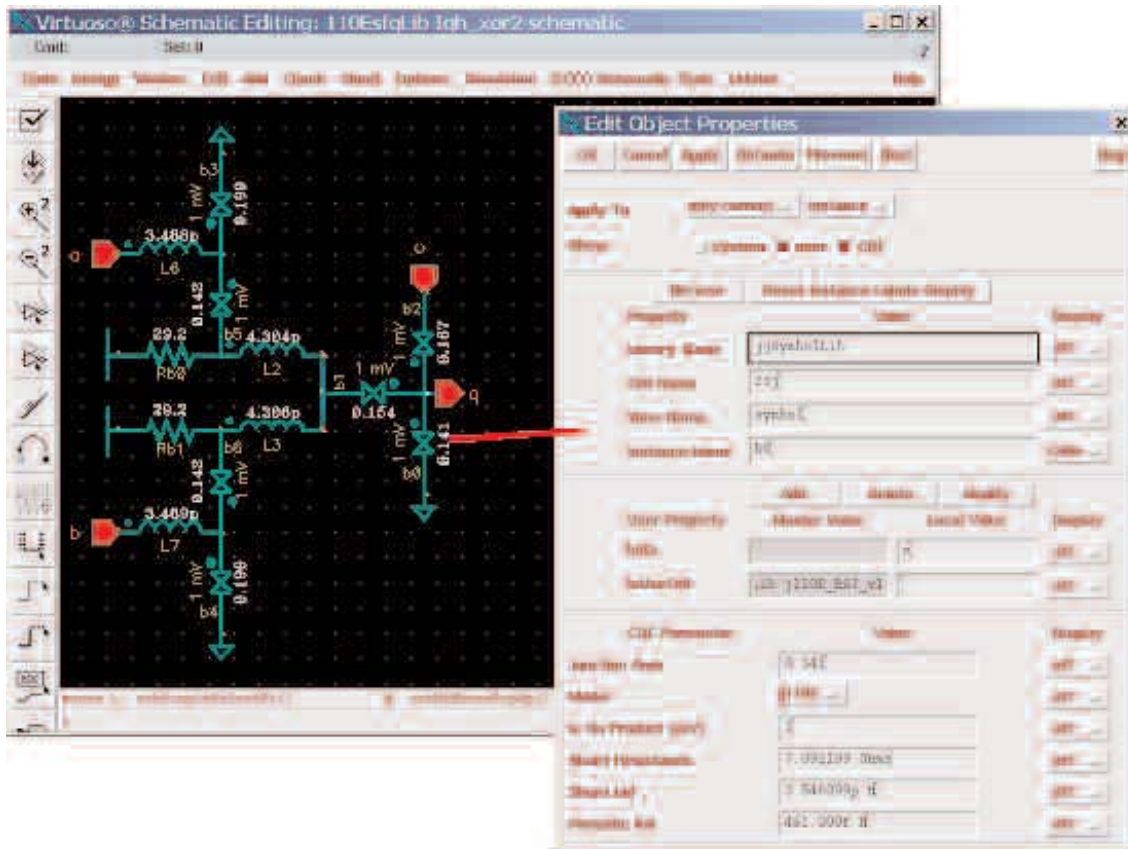


Fig. 2. The Schematic View describes the gate at the simplest level of individual devices and components.

The gate library, published by the foundry, consists of multiple cell views, including schematic, symbol, behavioral (VHDL in the present example), and physical layout. These will be described in turn, using a two-input XOR gate as an example. The schematic view is drawn by the library developer and describes the gate on the device level. As shown in Figure 2, the gate consists of resistors, inductors, and resistively-shunted Josephson junctions (JJs). The pop-up form is used to view and modify individual device parameters.




Fig. 3. The Symbol View is instantiated in a SPICE deck that describes physical-level simulation using WRSpipe. This is used by the gate library developer to optimize the gate and by the circuit designer to simulate circuits containing many gates.

The symbol view is used to place the cell in a larger circuit schematic. The schematic shown in Figure 3 is a SPICE deck that includes standard Josephson transmission line (JTL) input and output loads and input waveforms. This schematic is netlisted hierarchically down to the device level. This is used by the gate library developer to simulate the gate using WRSpipe, which contains the JJ device element. Physical-level parameter optimization is done using a combination of rule-based and software-based automated parameter selection. Larger SPICE decks containing many gate instances may be constructed by the circuit designer, netlisted to WRSpipe, and simulated.

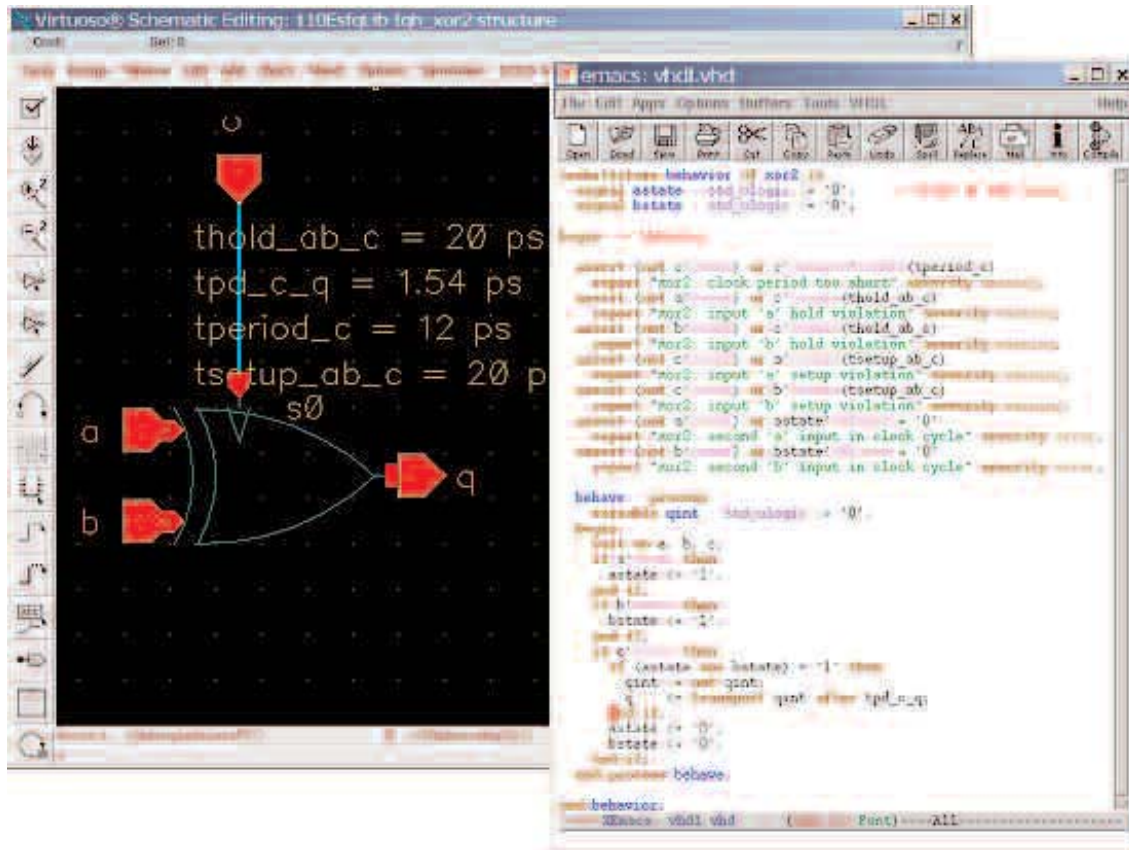


Fig. 4. The VHDL View captures the behavioral model of the gate in a hardware description language.

Each cell in the library also contains a VHDL behavioral model. As shown in Figure 4, the model follows a template containing:

- Internal state.
- Logical and timing violations.
- Logical operation.

Hierarchical circuit schematics of nearly arbitrary size and complexity can be constructed graphically by the circuit designer to be netlisted to VHDL for behavioral modeling.

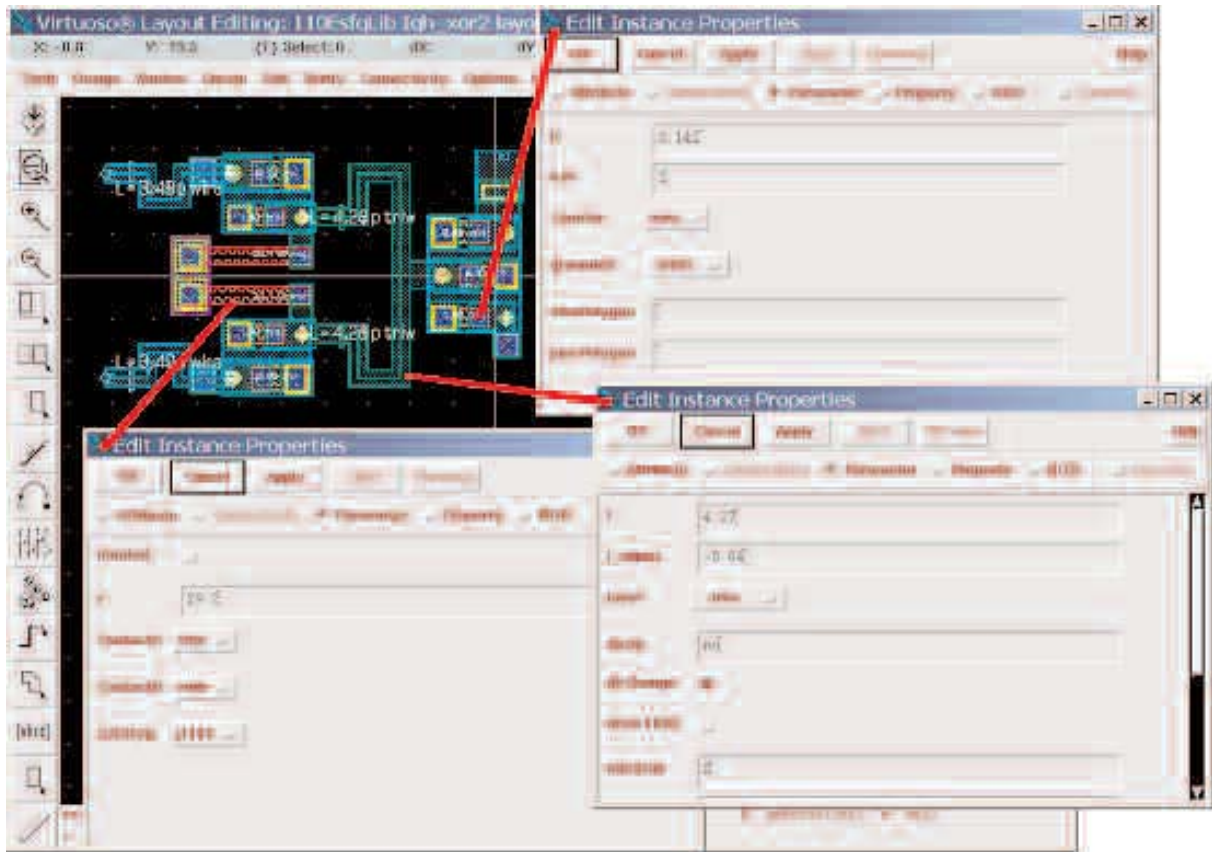


Fig. 5. The Physical Layout View uses parameterized cell (Pcell) components.

The physical layout view contains the mask data. This cell is constructed from parameterized cells (Pcells) for each element. In this regime, a common cell is instantiated for each kind of circuit element. Parameter values are then entered individually, and physical layout of the instance adjusts to reflect the entered values. Resistor, inductor and JJ Pcells and forms are shown in Figure 5. Physical layout is initially drawn by the library developer and can be copied out and modified by the circuit designer.

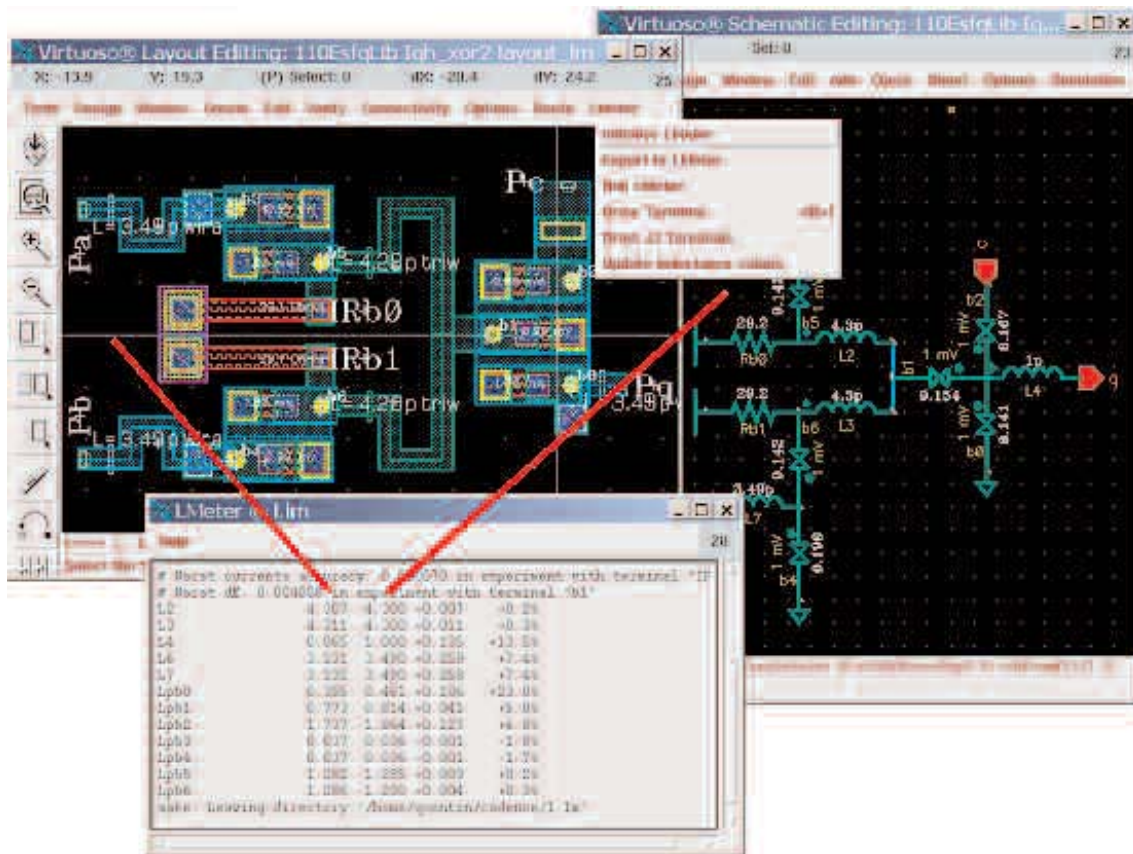


Fig. 6. LMeter is specialized software that extracts the inductance value of an interconnect from the physical layout. Inductance extraction is more important in RSFQ design than in semiconductor circuit design. Standard Cadence tools are generally inadequate for RSFQ.

Cadence supports extraction of device parameters from the physical layout. However, extraction of inductance values is not well-supported. Square-counting routines used in resistance extraction routines can be used, but are generally not accurate enough to be useful in typical layouts, particularly where parasitic inductance is involved. LMeter is custom software that extracts the inductance matrix. As shown in Figure 6, LMeter has a GUI incorporated into Cadence that can be used to compare layout values to the schematic and vice versa. This is typically used only by the gate library developer.

A significant challenge for fabrication at the 0.25-micron node is the development of inductance extraction software that is accurate for smaller features. The LMeter algorithm is 2D, i.e., it works from the assumption that metal lines are wide compared to the thickness of the dielectric. New software will use 3-dimensional modeling of magnetic fields to accurately predict the inductance of sub-micron lines.

We now turn to the subject of circuit design and verification, the area above the line in Figure 1. The challenge in the circuit design task is to manage the complexity of the design. Modeling of the circuit, which may be accomplished by any effective means, is more an art than a science. (Effective modeling is often distinct from mathematically accurate modeling.) For example, high-speed CMOS logic synthesis might be accomplished with simulation based on a hardware description language (HDL) and automated place-and-route. The HDL models need to contain the right information, such as gate delay as a function of the parasitic capacitive load associated with the physical wiring between gates. Standards for gate models exist, such as standard delay format (SDF) in which gate delay would be given in terms of nominal, typical high, and typical low values of physical parameters such as temperature.

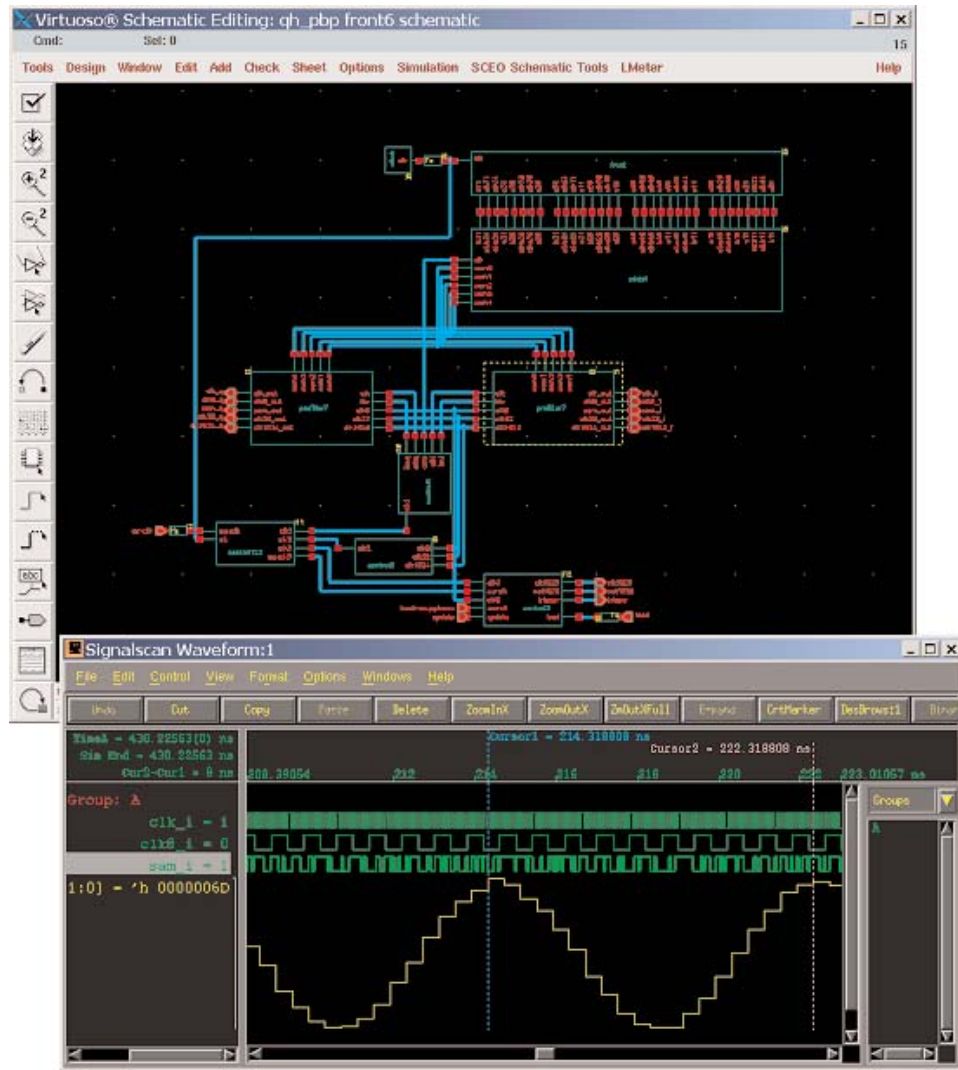


Fig. 7. VHDL simulation performed on a large, mixed signal superconductor circuit.

VHDL has been used in the design of complex multichip superconductor digital systems such as the programmable bandpass system that combined mixed signal, digital signal processing, and memory elements. Mixed signal VHDL simulation of the front-end chip is shown in Figure 7. The first-order design task is to insure that the circuit is logically correct. In this case VHDL is required to satisfy non-trivial requirements such as synchronization procedures between chips.

The second-order task is to insure correct signal timing, in this case using a simple but effective procedure that required data pulses to be time-centered between clock pulses at the input to each gate. This was implemented by setting hold and setup time of each gate to be slightly less than one-half the clock period.

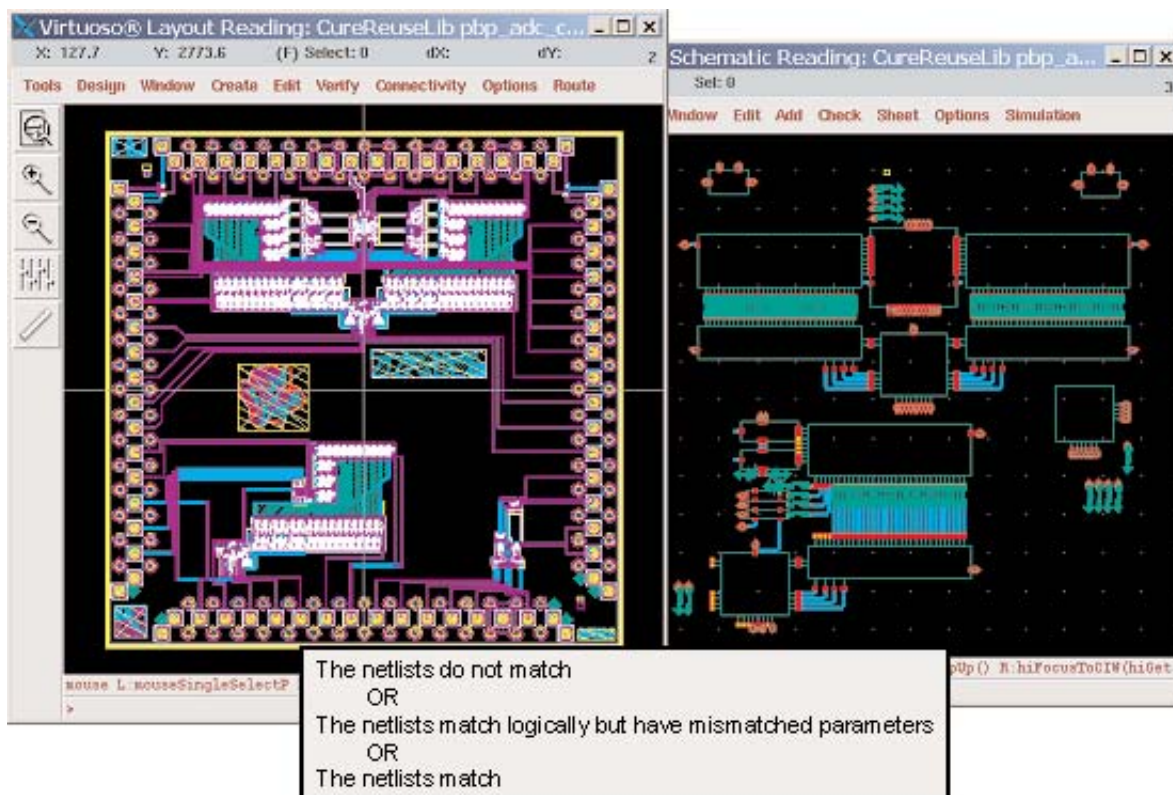


Fig. 8. Layout-versus-Schematic verification on the chip level for a large chip. LVS software checks the mask drawing against the circuit schematic, to identify errors such as open circuits, missing pieces, and extra components not included in the design simulation.

The same schematic that is used to design the circuit in VHDL can be used for layout-versus-schematic (LVS) verification. This is done all the way up to the chip level. Figure 8 shows an example chip containing a few thousand JJs that was verified in this way. Of course, design rule checking (DRC) is also performed on the entire chip.

First-pass success has become routine for chips of up to a few-thousand JJs manufactured in the Northrop Grumman foundry. This was achieved by the union of design verification and, on the foundry side, reduction of process control monitor and visual inspection data to characterize each die. Figure 9 is a gallery of chips, designed variously in 4kA/cm² Nb, 8kA cm² Nb, and 1kA cm² NbN foundry technologies that achieved first pass success. For these designs, functionality was achieved on the first die that was tested.

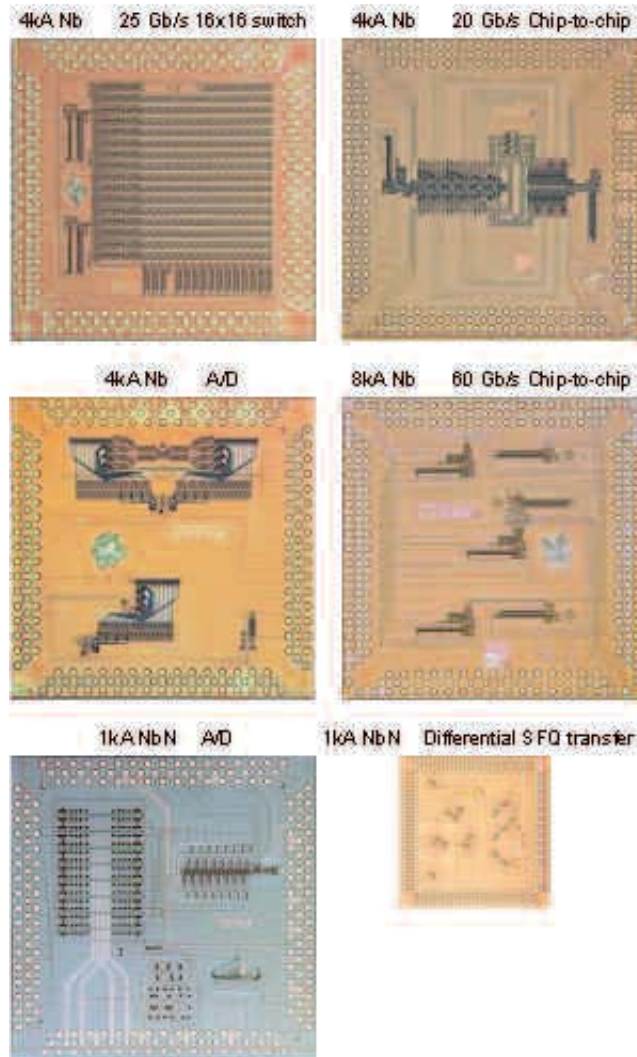


Fig. 9. First-pass success is routine for circuits of up to a few thousand JJs. Both 5mm chips and a 2.5 mm chip are shown.

Appendix K

DATA SIGNAL TRANSMISSION

ABSTRACT

This chapter evaluates the candidate technologies for data signal transmission from the core processors, operating at cryogenic temperatures, to the large (petabytes) shared data memory at room temperature and at some distance (tens of meters). The conclusion reached is that the only technology capable of providing the very large bandwidth requirements of the computer architecture required to meet government needs, the distances this data must be moved, and the operation of the superconductive processor at 4K, is massively parallel optical interconnects. Because of the limitations of current opto-electronic devices, a specific system configuration appears optimal: placing the 4K to 300K transmission components at an intermediate temperature within the cryogenic housing.

A number of alternative approaches are suggested whose goals are to:

- meet the exacting bandwidth requirements.
- reduce the power consumption of cryogenically located components.
- reduce size and cost by advances in component integration and packaging beyond the foreseeable needs of the marketplace in the 2010 time frame.

A listing of the specific efforts, their current status, and goals is shown, along with a roadmap to achieve the data signal transmission capabilities required to start the construction of a superconductive supercomputer by 2010.

Contents:

1. Current Status of Optical Interconnects
 - 1.1. Comparison of Optical and Electrical Interconnects for Superconductive Computers
 - 1.2. Current Optical Interconnect R&D Efforts
2. Readiness for major investment
3. Issues and concerns
 - 3.1. Data Rate requirements
 - 3.2. Issues of Using Optics at 4K
 - 3.3. The 4K-Intermediate Temperature Electrical Option
 - 3.4. Coarse vs. Dense Wave Division Multiplexing (CWDM vs. DWDM)
 - 3.5. Developments for Low Temperature Opto-electronic Components
4. Conclusions
5. Roadmap
6. Funding

1. CURRENT STATUS OF OPTICAL INTERCONNECTS IN COMPUTERS

The term *optical interconnect* generally refers to short reach (< 600 m) optical links in many parallel optical channels. Due to the short distances involved, multimode optical fibers or optical waveguides are commonly used. Optical interconnects are commercially available today in module form for link lengths up to 600 m and data rates per channel of 2.5 Gbps. These modules mount directly to a printed circuit board to make electrical connection to the integrated circuits and use multimode optical ribbon fiber to make optical connection from a transmitter module to a receiver module. Due to the increase in individual line rates, the increase in the aggregate data rate for systems, and the need to increase bandwidth density, there is a need to move optical interconnects closer to the I/O pin electronics. This will require advances in packaging, thermal management and waveguide technology, all of which will reduce size and costs. For this reason, this area is the subject of intense study worldwide, with some efforts funded by industry, and others by governmental entities such as DARPA. All of these efforts are aimed at achieving devices and systems that will be practical for use in the near future for conventional semiconductor based computers, in terms of cost and power consumption as well as performance.

1.1. COMPARISON OF OPTICAL AND ELECTRICAL INTERCONNECTS FOR SUPERCONDUCTIVE COMPUTERS

The necessity for optical interconnects for backplane and inter-cabinet high bandwidth connections is best demonstrated by Figure 1 which shows simulated data of frequency-dependent loss for a 50 cm long electrical interconnect on conventional circuit boards. Above 1 GHz, dielectric losses rapidly become larger than skin effect losses. The result includes the effect of packages, pad capacitance, via inductance, and connectors, in addition to the loss associated with the traces in FR4 circuit board material. Newer material can extend the bandwidth of electrical interconnects by over two times. While frequency-dependent loss is considered the primary problem, the effects of reflections and crosstalk on the performance of electrical interconnect can pose other challenges for interconnect designers as the data rate increases beyond 10 GHz. In addition, there is literature that compares the power consumption, at room temperature, of electrical and optical transmission which establishes that for data rates higher than 6-8 Gbps, the distance that electrical transmission is advantageous between computer elements does not exceed 1 meter, and is frequently much less. For short distances, especially inside the 4K (superconductive) environment, however, electrical connections are to be preferred for their simplicity and low power.




Figure 1. Channel Loss for 50 cm Electrical Link (Intel Technology Jnl, 8, 2004, p116)

Some studies, such as those at Intel, and a number of universities, are aimed at ultra-short reach applications, such as chip to chip and on-board communications. These are of little interest for a superconductive computer. The major advantage of optics is in the regime of backplane or inter-cabinet interconnects where the low attenuation, high bandwidth, and small form factor of optical fibers dominate; in the case of a superconductive processor, the thermal advantages of glass vs. copper are also very important.

The large scale of a superconductive based petaflop class computer plays a major role in the choice of interconnect technology. Since a petaflop computer is a very large machine, on the scale of tens of meters, the interconnects between the various elements will likely be optical. Furthermore, the number of individual data streams will be very large; one particular architecture, shown in Figure 2, would require $2 \times 64 \times 4096 = 524,288$ data streams at 50 Gbps each between the superconductive processors and the SRAM cluster alone, and perhaps many times larger between higher levels. Only the small form factor of optical fibers along with the possibility of using optical Wavelength Division Multiplexing (WDM) is compatible with this large numbers of interconnections. The very low thermal conductivity of the 0.005" diameter glass fibers, compared to that of copper cables necessary to carry the same bandwidth represents a major reduction on the heat load at 4K, thereby reducing the wallplug power of the system substantially.




Figure 2. Long path data signal transmission requirements for one proposed petaflop architecture. This would require over 500,000 data channels at 50 Gbps each, assuming 64 bit words.

1.2. CURRENT OPTICAL INTERCONNECT R&D EFFORTS

Optical interconnect technology for distances of <300 meters has been the subject of several research efforts recently, much of which is being sponsored by DARPA. A major packaging issue, the necessity to precisely align fiber to laser, is greatly eased in these short range systems through the use of multimode fiber to transmit the data. A number of industrial laboratories have been involved, including Agilent, IBM, Intel, Sun, Mayo, and Infineon, as well as major universities such as MIT, Stanford, U. of Arizona, UCLA, UCSB, U. of Texas, U. of Colorado, Ruhr U. (Germany) and the U. of Bristol (UK). For example, using Vertical Cavity Surface Emitting Lasers (VCSELs) and Coarse WDM the joint IBM/Agilent effort has achieved 240 Gbps aggregate over 12 fibers, each carrying four wavelengths at 5 Gbps each. The next step is planned to be 480 Gbps, with each channel at 10 Gbps. The current stage of this effort is illustrated in Figure 3.

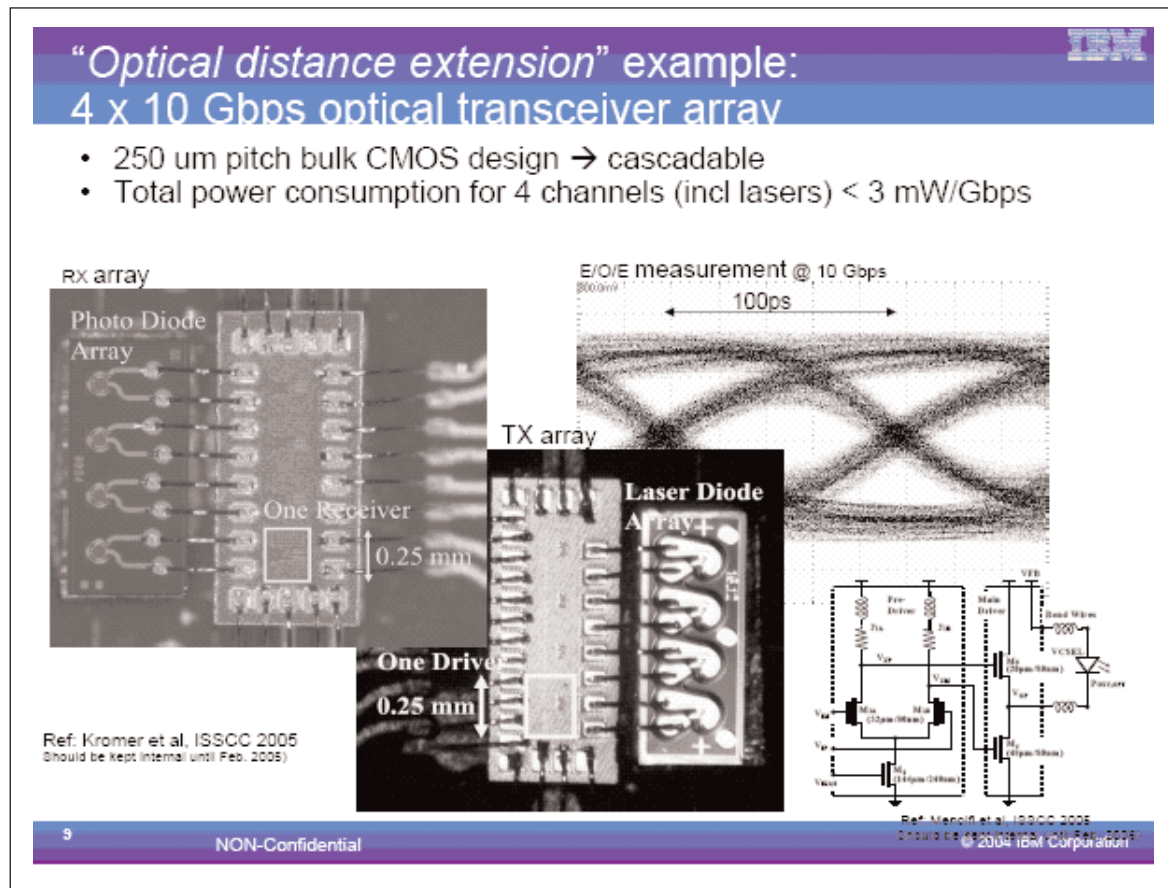


Figure 3. Four channel transceiver arrays produced by the IBM/Agilent effort funded by DARPA

Work at the U. of Colorado has a goal of building an array of 40 Gbps directly modulated VCSELs; they have already achieved 10 Gbps, and expect to reach 20 Gbps in the near future. The U. of Texas at Austin is working on amplifierless receivers which are crucial for 4K applications, while the Mayo effort is focused on Bit Error Rate characterization and test-bed applications. Much of the Intel effort is aimed at achieving ultra-short reach applications for chip-to-chip interconnects. All of these efforts are aimed at achieving devices and systems that will be practical for use in the near future in terms of cost and power consumption as well as performance.

A development in long haul communications technology which is rapidly approaching implementation into 40 Gbps systems is the use of differential phase shift keying (DPSK) in place of the conventional on-off keying normally used (and used in the above described systems). While the receivers are slightly more complex, this allows the use of simpler transmitters, which is of great advantage in this application (see Sections 3.2 and 3.3)

There are also a number of very interesting developments in exotic materials and devices which have the potential to greatly reduce the cost of these components and systems well beyond the scope of the efforts mentioned above. One area, funded under the DARPA MORPH program, includes the development of organic glasses and polymers which may be capable of producing optical modulators with very low modulation powers, as well as allowing simple fabrication techniques for large arrays. Another area, once funded by ONR, is the use of magneto-optic modulators directly driven by the superconductive circuitry. This would be the optimal solution since it would eliminate the need for any amplifiers and would operate at 4K. While in an earlier stage of development than the abovementioned efforts, one of these may become available in time for the construction of a full scale petaflop supercomputer starting in 2012, while a demonstration system could be built using the less advantageous technology available in 2010.

2. READINESS FOR MAJOR INVESTMENT

There is considerable interest in some parts of the computer industry to move into optical interconnects, as is evidenced by the IBM and Agilent co-funding of their above described efforts sponsored by DARPA under the Terabus program, which are described above. DARPA funding has been about \$7.5M over an 18 month period. It can be expected that the companies and universities who are already involved in this effort could absorb a major increase to speed up the effort without straining their resources. In fact there are a number of other companies who have already expressed interest in this area.

Though Intel interests are broader than only short range (<300m) interconnects, the technologies they are developing for packaging telecommunications data rate components will help to reduce costs and form factors. Though their current focus is on 10 Gbps components, the use of 40 Gbps line rates in telecommunications is increasing. It would be most helpful in the superconductive supercomputing context if this increased throughput was achieved by increasing the data rate per color rather than by increasing the number of colors used in the WDM scheme. Their interest in low cost DPSK receivers would be significant, as would that of router manufacturers such as Cisco.

Some technical issues must be resolved (see section 3) for the use of optical interconnects in a superconductive computer system which are not currently being addressed. Some of these are performance related, some are environmental, and some economic; however almost all are in the category of "not if, but when". A major investment could greatly speed the development of the key technologies to provide them at the time when they will be needed.

3. ISSUES AND CONCERNS

3.1 DATA RATE REQUIREMENTS

The use of superconductive processors implies that very high clock speeds will be used, which makes it very desirable to use the commensurate data rates for the interconnects. The following discussions assume a processor clock speed of 50 GHz, and a data rate for an optical interconnect of 50 Gbps. This is not unreasonable given the commercial use of 40 Gbps data rates (OC-768), frequently associated with forward error correction, raising the total transmission rate to 43+ Gbps. Commercially driven telecommunications technology is already generating optical and electronic components capable of operating at such rates, although as yet these components are expensive. While a number of laboratories have produced serial data rates of >100 Gbps, the widespread use of such rates is not likely in the next 5-10 years; therefore production technology to meet the 2010-2014 target dates for a superconductive supercomputer will not be available.

If technology, power, or cost limitations of superconductive clock speed optical technologies prove to be problematic, one possible option which should be held open would be to go to half clock rate transmission speeds. While this necessitates 2 bit buffer stages on the superconductive chip, this should not present any serious circuit design issues. If this approach were taken, 20-25 Gbps opto-electronic components could be used, which should be much more readily achievable within the target time frame of 2010 for a demonstration system.

It is assumed throughout this chapter that the word width is 64 bits; the use of a 72 bit wide word so that error correction lines can be included in the data transmission is not an issue.

3.2 ISSUES OF USING OPTICS AT 4K

Given the use of very low power consumption of superconductive processors, the total power expended by the inter-level communications becomes a major issue in the system. With the power efficiency of refrigeration systems at 4K at 0.2%, because of the power demands, placing *transmitting* optics at 4K does not appear to be feasible, given the options considered below.

a) Placing optical sources at 4K is very power consuming. This is not only due to the power needed to drive one laser/line (x64 lines wide) but also for the wide bandwidth analog amplifiers required to raise the voltage levels from the (typically) 3 mV characteristic of Josephson Junction (JJ) technology to the 1.5 -2 volts required to drive the lasers. There are good fundamental reasons to expect that this will not change in the near future. The power goals of the current DARPA efforts are in the 3-10 mW/Gbps range for 10-20 Gbps devices.

b) Another option is to generate CW laser light at room temperature and transmit it by fiber into the 4K environment where it would be modulated by the data on each line and then transmitted on to the next level, located at room temperature. This would be a very attractive option if an optical modulator capable of operating at 50 Gbps, with a drive voltage (3-5 mV) consistent with superconductive technology existed. Currently available devices at 40 GHz require drive voltages in the range of 6 volts p-p, with rf drive powers of 700 mW into a 50 ohm line. Furthermore this would also require the use of high gain, large bandwidth amplifiers, not consistent with very low power consumption. There are modulator devices being considered which might have the desired properties. These include the coupling of micro-ring resonators to standard Mach-Zender modulators to reduce the modulation voltages required to less than one volt, and the development of micro-disk electro-absorption modulators operating at 40 Gbps. A third possibility is the use of VCSELs as modulators. In this case they act as injection locked lasers which have a higher modulation frequency capability than normal directly modulated lasers. The longer term possibilities include the use of organic glasses or polymers as either amplitude or phase modulators, which have much lower drive power requirements. or the use of magneto-optic modulators directly driven by the superconductive circuitry which would eliminate the need for any amplifiers. All of these alternatives are in a very early stage of research, and may not be available in time to meet the needs of the STA, with the currently planned funding rate. These options should be evaluated as they develop.

c) Even if the desirable modulators discussed in section 3.2(b) were available, they may solve only half the problem. An optical receiver would also have to be placed at 4K to capture the data flow into the processor. Reasonable detected optical powers (200μw) would generate adequate voltage (100-160 mV) to drive the superconductive circuitry directly if a 100 ohm load and appropriate thresholding can be done.

d) A simpler and more elegant solution is the use of amplifierless receivers. This is being explored for use in Si-CMOS devices by U.Texas at Austin. Since superconductive circuits are current driven, the superconductive logic could be driven directly with the detected current (100-160 μA) from the photodiodes. If this is achievable, it may be very advantageous to have the 300K to 4K interconnect be all optical, even if the reverse path is not. However, if any electronic amplification were required, the amplifier power would likely become a major thermal problem. This issue must be closely examined.

3.3 THE 4K-INTERMEDIATE TEMPERATURE ELECTRICAL OPTION

Since placing some optical components at 4K does not appear to be feasible, other approaches must be considered. The design of the cryogenics appears to be compatible with having a short (<5 cm) section of high speed flexible ribbon cable, described elsewhere in the Systems Integration section of this report, connecting the 4K section of the system with a section at some intermediate temperature, say 40+20K. This transmission line would carry 64/128 copper strip lines (perhaps bidirectional 64 bit word) at a pitch of 100 lines/inch. Because of its small size and low loss (perhaps 3 dB) it would be capable of carrying data between the 4K and intermediate temperature environments without a severe thermal penalty. The power efficiency of 40K refrigeration is ~2%, making the use of optical transmitter and receiver components more feasible from a power standpoint. The issue of the incompatibility of voltage levels between superconductive circuitry and optics can now be addressed using advanced technology amplifiers, such as InP HEMT transistors, which have been shown to operate at low temperatures with adequate bandwidth. It should be pointed out here that since high gain linear amplifiers tend to be power hungry and the data is digital, linear amplification need only be used to provide enough gain to reach a level (~300 mV) at which a (less power hungry) threshold circuit can operate.

The bandwidth capability of electrical ribbon cables are at present uncertain. They are currently used up to 3 Gbps, and studies indicate that they are useful up to 20-25 Gbps and perhaps higher. This could necessitate running data lines to intermediate temperature at half of the processor clock speed. In this case, the superconductive processor would have to provide twice the number of half speed lines. These could be re-clocked to full data rate at the intermediate temperature, or else double the number of optical lines would have to go to room temperature. Either is likely to be feasible, and is a system design issue to be determined when the individual components have been further developed. Another possibility which might be investigated is the use of high temperature superconductors (HTS) in place of copper for the transmission lines, since the upper temperature is compatible with known HTS materials. Some research has been done in this area, but the feasibility for use of these materials on flexible cables has not been fully developed. These issues are addressed in more detail in the Systems Integration chapter.

If this option is taken, the receiver can be placed at either 4K or the intermediate temperature, and the issue will likely be determined by detailed engineering considerations rather than more fundamental concerns. The data transmitter will be at the intermediate temperature, but the option of where the light is generated (at the intermediate temperature or 300K) remains an open issue. If suitably low powered modulators can be developed in time, it might be thermally advantageous to generate CW light at 300K. Otherwise very efficient VCSELs at the intermediate temperature might prove more advantageous. Again, an engineering trade-off, to include component reliability must be made midway in the component development process. VCSELs, as active devices, may be less reliable than passive modulators.

An issue which must be thoroughly explored is the use of room temperature optical components in a cryogenic environment. Optical fiber has routinely been used at low temperatures in laboratory and space environments. Also, 3-5 semiconductor components and many electro-optic materials have performance characteristics at low temperatures superior to 300K operation. However telecommunications grade optical components, such as fiber ribbon cables, connectors, wave division multiplexers, etc must be evaluated for use at low temperatures. This should be done early in the program to ensure that appropriately engineered components are available when needed.

Further developments in all these areas might be needed. All the components necessary for this option must be evaluated in detail for a system, but it is not likely to become a major issue.

3.4 COARSE VS. DENSE WAVE DIVISION MULTIPLEXING (CWDM VS. DWDM)

The use of WDM for optical interconnects has been suggested by a number of authors, and demonstrated recently by Agilent under the DARPA Terabus program. This effort uses 4 wavelengths of Coarse WDM at 30 nm (5.5 THz) spacing from 990 to 1080 nm. The optical multiplexing components used here are compatible with the use of multimode fiber, which is highly advantageous for a system with a large fiber count from the standpoints of light collection efficiency and ease of alignment. However, it should be pointed out here that this may not be compatible with 50 Gbps operation, where the detector size may become smaller than the multimode fiber core size, entailing large optical coupling losses. This must be evaluated in a complete system design study. To achieve 64 channels (a word width) on a single fiber would require the use of the closer wavelength spacings, more typical of the standard Dense WDM spacing of 0.8 nm (100 GHz at 1550 nm). In this case it can be anticipated that single mode fiber and optical multiplexing components would be required, with the associated losses and tighter alignment tolerances. However, only a single fiber would be required for a word width, a considerable advantage for a large system. In addition, the use of single mode fibers would enable the use of Differential Phase Shift Keyed (DPSK) modulation, at least for the Intermediate Temperature to 300K link. This could lead to simpler modulators and improved performance.

Another issue involved in the choice between CWDM and DWDM is that of manufacturability. This is more easily achieved in CWDM, where say 4 arrays (each array at a different wavelength) of 16 lasers each can be placed side by side to achieve 64 channels, than for DWDM where an array of 64 VCSELS with each laser operating at a different wavelength is required. A conservative data transmission concept using CWDM operating at half clock rate (25 Gbps) is illustrated in Figure 4. Given the successful achievement of the DARPA TERABUS goals by 2007 this technology should be commercially available by 2010.

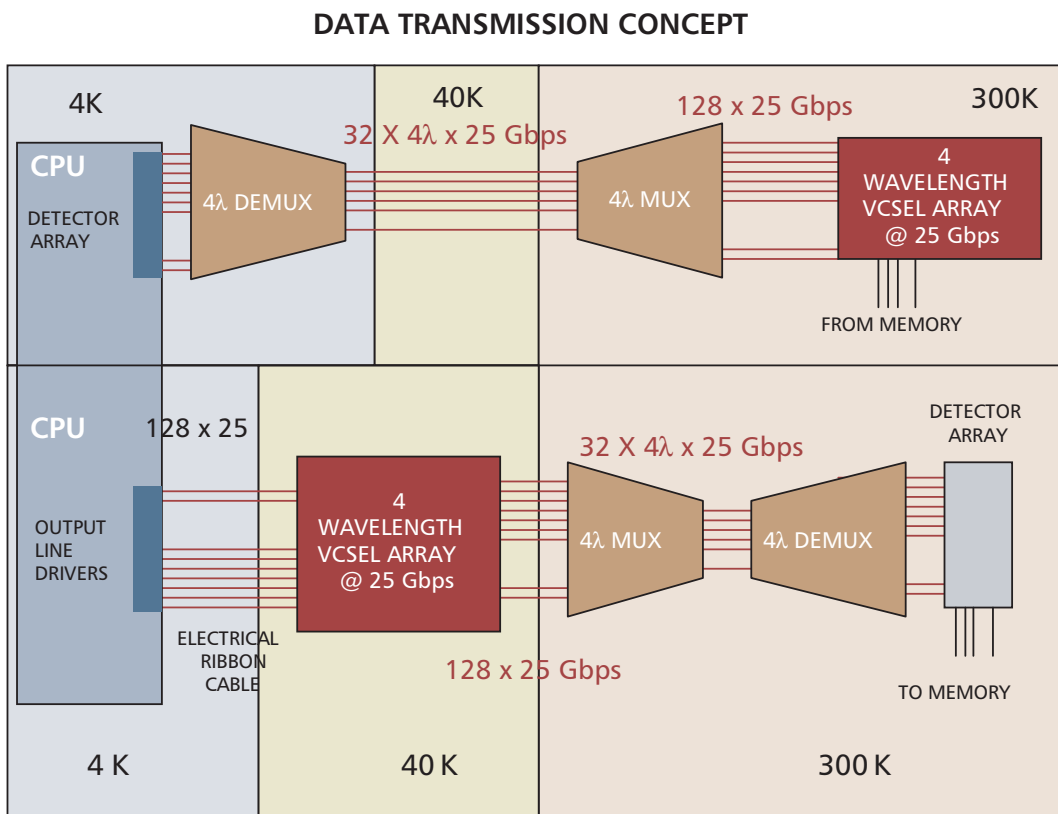


Figure 4: A 64-fiber, 4-wavelength, 25-Gbps CWDM System for bidirectional transmission totaling 6.4 Tbps between a superconductive processor at 4 K and high speed mass memory at 300K. Optical connections are shown in red, electrical in black. This technology should be commercially available by 2010.

There is another option to employ DWDM which might be considered however, if, as discussed above, it is possible to make an electrically and optically efficient 50 Gbps modulator capable of operating at the intermediate temperature. In this case, a light source placed at 300K that is a laser whose output is not a single wavelength, but a comb of wavelengths, could be used. Such lasers have been built for research purposes and are now commercially available from at least one vendor. In this case, the laser light is transmitted by a single fiber to the intermediate temperature, where it is demultiplexed to an array of *identical* modulators, and remultiplexed onto a single output fiber to 300K. This technique would eliminate the need to control the wavelength of very large numbers of individual lasers. There would likely be a major wallplug power advantage here since the electrical power to modulate the light at 40K would then be less than the power needed to generate it there. Figure 5 shows a schematic representation of a DWDM system which employs only 3 fibers between cryogenic components and room temperature.

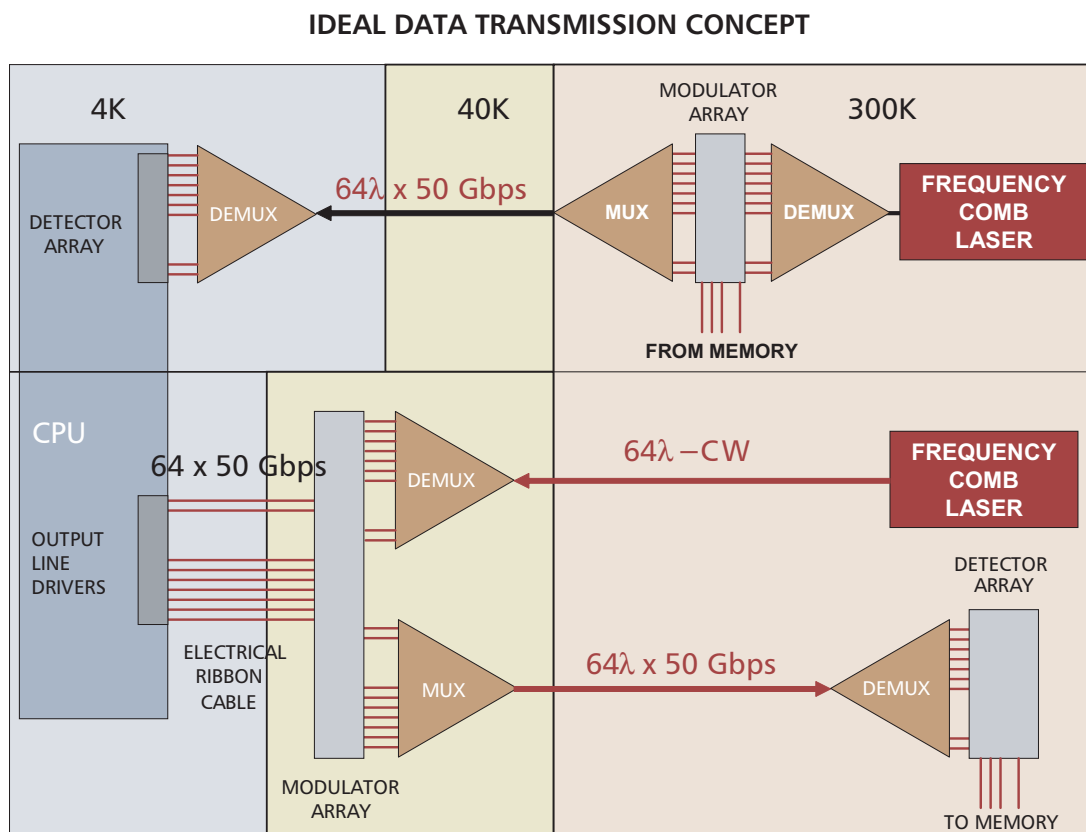


Figure 5. A 3-fiber, 64-wavelength, 50-Gbps DWDM System for bidirectional transmission totaling 6.4 Tbps between a superconductive processor at 4 K and high speed mass memory at 300K. Optical connections are shown in red, electrical in black.

To accomplish at least one of these options by 2010 appears to be feasible, given the current state of the technology and the estimated cost to achieve the desired goals by then. The development of more easily manufacturable technologies, such as organic glasses or polymers, for application by 2012 should receive special attention because of cost issues. As discussed earlier, the final choice will be determined as much by systems engineering and economic decisions as by the availability of any of the required technologies. Each of the options should be examined in detail in the initial phase of the study, with a commitment to the final development of a manufacturing capability to follow.

3.5 DEVELOPMENTS FOR LOW TEMPERATURE OPTO-ELECTRONIC COMPONENTS

With the exception of InP HEMT amplifiers, which have been used at low temperatures, optical and opto-electronic components which are made to operate at an intermediate temperature are generally not available off-the-shelf. In addition, all the current development efforts are aimed at room temperature use. There do not appear to be any fundamental issues which would preclude the necessary developments; the initial phases of the effort should focus on identifying those areas which require adaptation of standard room temperature optical and opto-electronic technologies for use at low temperatures.

4. CONCLUSIONS

Given the efforts outlined briefly above, optical interconnect technology capable of handling the 50 Gbps data rates and overall data transmission requirements for a superconductive petaflop computer can be made available by 2010. Commercially available telecommunications technology is now at 40-43 Gbps per single data stream; commercial off-the-shelf optical interconnect technology is currently at 2.5 Gbps. Efforts to increase these rates substantially, and reduce cost, size and power consumption are ongoing under a DARPA program with company cost sharing. These programs will provide advances in the near future; however, the more advanced and specialized requirements of a superconductive supercomputer will require additional government funding to meet the 2010 time line, as well as to provide components suitable for the cryogenic environment in which many of them must operate.

The key issues to be addressed are the following:

- Developing device structures capable of 50 Gbps operation in highly integrable formats.
- Reduction of electrical power requirements of large numbers of electro-optical components operating at cryogenic temperatures.
- Packaging and integration of these components to reduce the size and ease the assembly of a system of the scale of a petaflops supercomputer.
- Reducing the cost of these components by large scale integration and simple and rapid manufacturing techniques.

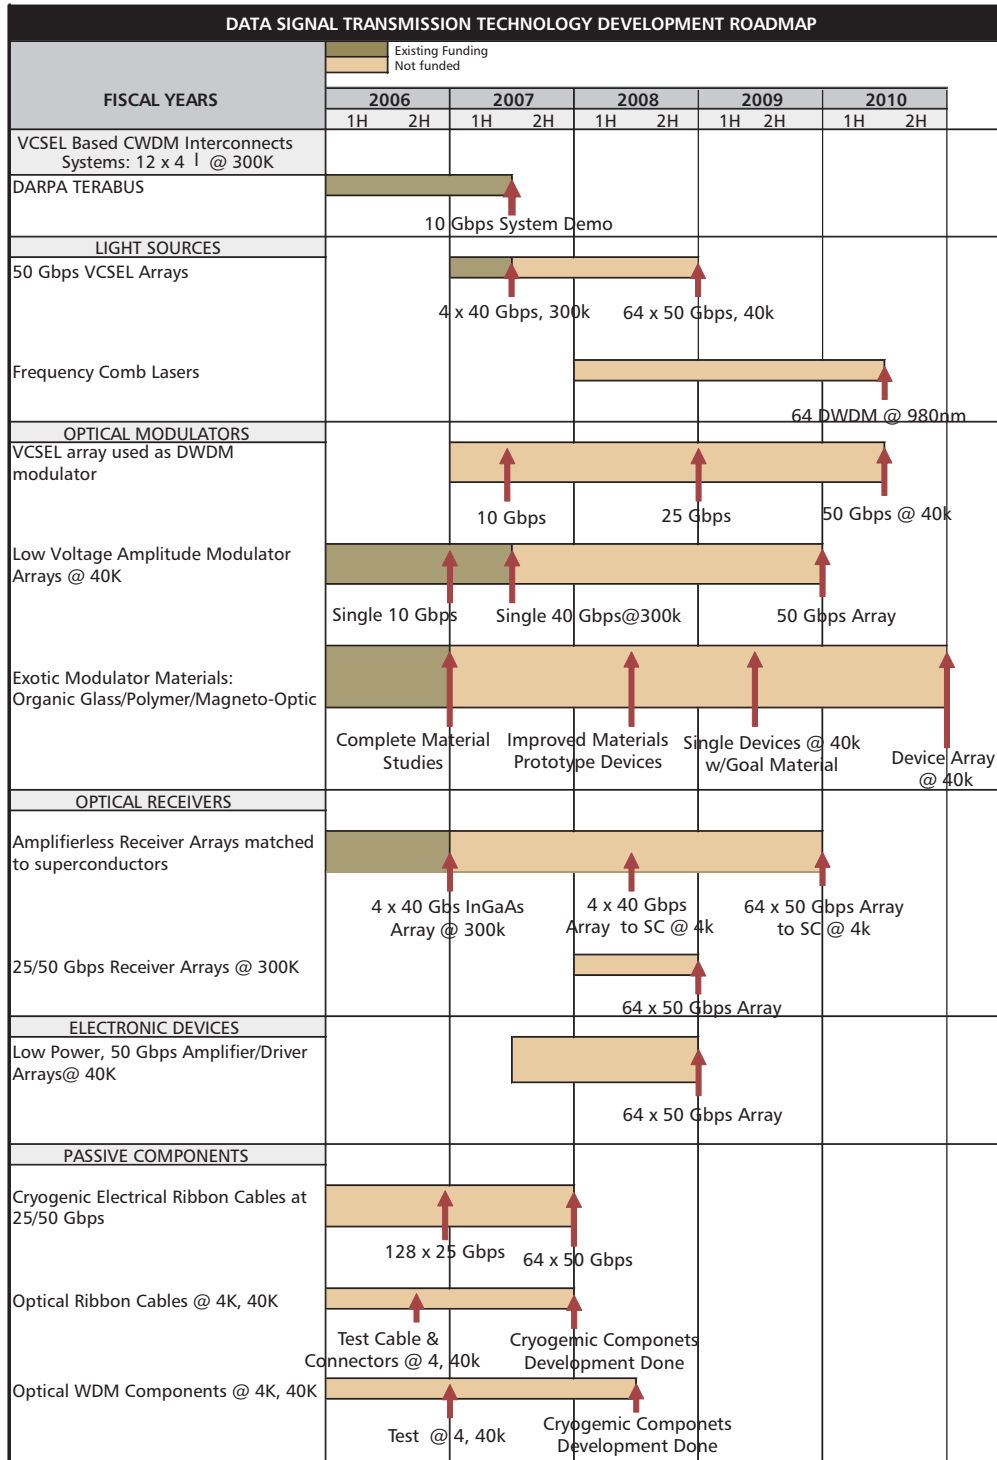
Table 1 is a summary of the developments required to fully develop the optical interconnect components needed to begin construction of a superconductive supercomputer by 2010.

TABLE 1. COMPONENT TECHNOLOGY DEVELOPMENT SUMMARY

DEVICE	CURRENT STATUS 2005	DESIRED STATE 2010	CRYOGENIC USE TESTED	ONGOING R&D
VCSEL Based CWDM Interconnect Systems: 12x4λ @ 300K				
DARPA Terabus	12 x 4λ x 10 Gbps	8 x 8λ x 50 Gbps	N	Y
LIGHT SOURCES				
50 Gbps VCSEL Arrays	10 Gbps	50 Gbps	Y	Y
Frequency Comb Lasers	32λ x 1550 nm	64λ x 980 nm	N/A	N
OPTICAL MODULATORS				
VCSEL used as DWDM modulator	2.5 Gbps	50 Gbps	N	Y
Low Voltage Amplitude Modulator @ 40 K	4-6 volts On-Off/ 40 Gbps Single Devices/ COTS	0.5 volt on-off/ 50 Gbps Arrays	N	Y
Exotic Modulator Materials: Organic Glass/Polymer/Magneto-Optic	Materials studies show 3x improvement over conventional materials	6-10 fold improvement in devices	N	Y
OPTICAL RECEIVERS				
Amplifierless Receivers matched to superconductors	20 Gbps demonstrated	50 Gbps Array	N	N
25/50 Gbps Receiver Arrays @ 300K	50 Gbps Single Devices – COTS 4 x 40 Gbps arrays Lab	Ideally, word-wide arrays	N/A	N
ELECTRONIC DEVICES				
Low Power, 50 Gbps Amplifier/Driver Arrays @ 40 K	Single Amplifiers (InP-HEMT) – COTS	Lower power, ideally, word-wide arrays	Y	N
PASSIVE COMPONENTS				
Cryogenic Electrical Ribbon Cables at 25/50 Gbps	3 Gbps	25/50 Gbps	Y	N
Optical Ribbon Cables @ 4 K, 40 K	COTS meet Commercial Specs/ Fibers OK at 4 K	Cables at 4 K	N	N
Optical WDM Components @ 4 K, 40 K	COTS meet Commercial Specs	4 K Operation	N	N

5. ROADMAP

The roadmap below shows the specification and the proposed schedule of the critical project stages related to the development of the interconnect technology from cryogenic processors to room temperature mass memory systems for a system demonstration. The roadmap assumes full-government funding. There are no known industrial companies in the U.S. that would fully support this work at this stage.



6. FUNDING

There are efforts that will have to be funded specifically to meet the technical requirements for superconductive supercomputers, especially for low temperature operation (4K/40K) and perhaps increasing the serial data rates from the 40-43 Gbps telecommunications standard to the 50+ Gbps superconductive processor clock speed. A total of \$71M should insure that the necessary components will be available by 2010. This does not include the amount necessary to develop the electrical ribbon cable interconnect between the 4K and intermediate temperature environment, which is included in the Systems Integration Chapter.

Appendix L

MULTI-CHIP MODULES AND BOARDS

In order to accommodate thousands of processor and other support chips for a SCE-based petaflops scale Supercomputer, multi-chip packaging that minimizes interconnects latency (i.e., Time of flight (TOF) and minimizes the total area and volume is needed. A well designed hierarchy of secondary packaging starting from SCE and memory chips and including multi-chip modules (very dense substrates supporting these chips) and printed circuit boards housing MCMs is needed.

The secondary packaging problem for large-scale SCE systems was investigated in previous programs such as HTMT¹. A typical approach is shown in Figure 1. This package occupies about 1 m³ with a power density of 1 kW at 4 K. Chips are mounted on 512 multi-chip modules (MCM) that allow a chip-to-chip bandwidth of 32 Gbps per channel. Bisectonal bandwidth into and out of the cryostat is 32 Pbps.




Figure 1. Packaging concept for HTMT SCP, showing 512 fully integrated multi-chip modules (MCMs) connected to 160 octagonal printed circuit boards (PCBs). The MCMs, stacked four high in blocks of 16, are connected to each other vertically with the use of short cables, and to room temperature electronics with flexible ribbon cables (I/Os). The drawing has one set of the eight MCM stacks missing, and only shows one of the eight sets of I/O cable.

Each MCM, up to 20 cm on a side, has up to 50 SCE chips and 8 processor units. Flip-chip bonding is used to attach the chips to the MCM providing high interconnect density (2,000-5,000 pinouts per chip), multi-Gbps data rates, automated assembly, and reworkability. The cutoff frequency of channels between chips is in the THz regime. The vertical MCMs are edge-mounted to 160 horizontal multi-layer printed circuit boards (PCB) which allows for MCM-to-MCM connections. Adjacent MCMs are connected to each other along their top and bottom edges with flexible ribbon cable. These cables consist of lithographically defined signal lines on a substrate of polyimide film in a stripline configuration, which have been demonstrated in a cryogenic test station at data rates up to 3 Gbps. A cryogenic network (CNET) allows any processor unit to communicate with any other with less than 20 ns latency.

¹ HTMT Program Phase III Final Report, 2002

The bandwidth, chip density and interconnect distance requirements place multi-chip modules (MCM) as the leading candidate technology for large SCE-based systems. The MCM requirements include:

- Use of niobium wiring for the dielectric portion of the MCM (MCM-D).
- A combination ceramic and dielectric MCM (MCM C+D) technology.
- Low impedance transmission lines.

The MCM is projected can hold 50 chips on a 20 cm X 20 cm substrate and pass data at 32 Gbps (Figure 2). It will have multiple wiring layers (up to nine superconductor wires), separated by ground planes and will need to accommodate different impedance wiring. To date, 20 SCE chips have been attached to an MCM of this type, each chip with 1,360 bumps.




Figure 2. A multi-chip module with SCE chips(left, NGST's Switch chip MCM with amplifier chip; center, NGST's MCM; right, HYPRES' MCM).

Although MCMs with SCE chips are feasible, the signal interface and data links impose further challenges. The large number of signal channels and high channel densities is difficult to achieve for two reasons:

1. The signals pass from the chips into the MCM layers within the area constraints of the pad geometry (maximum connection density for this process is dominated by via diameter and via capture pad area).
2. The MCM must transport signals laterally between chips and to peripheral connectors (maximum density for lateral signal transmission is determined by the wiring pitch and the number of layers).

1) Chip-to-MCM I/O: The issues for the die attach are the pin-out density (2 k to 5 k signal pads per chip), and the inductance (< 10 pH) and bandwidth requirements (> 50 Gbps). The leading candidate for the chip-to-MCM attach is flip-chip bonding utilizing solder reflow connections. This attach method allows for both high interconnect densities as well as the low-inductance connections required for 32 Gbps data transmission. Figure 3 gives interconnect inductance as a function of bond length or height for different types of interconnects. Other candidate technologies typically suffer from either insufficient interconnect densities or relatively large inductances which limit data rates to less than 10 Gbps.

A data rate of 32 Gbps between chips appears feasible using appropriate superconductor driver and receiver circuits. Bandwidth is limited by a L/R time constant, where L is the inductance of the chip-to-MCM bump and R is the impedance of the microstrip wiring. For reasonable values of L (3-10 pH) and R (10-20 Ω), the cutoff frequency is in the THz regime.




Figure 3. Lead inductance vs. bond length or height for different attach techniques.

Use of RSFQ logic circuits offers the advantage of ultra-low power, but the signal levels are low, so chip-to-chip bandwidth is limited unless on-chip driver circuits are used to boost the signal level. Such off-chip drivers need to match the impedance of the driver circuits and the transmission line interconnects. To date, 10 Gbps asynchronous off-chip drivers have been demonstrated. Alternatively, simple synchronous latches may be used. With the planned reduction in feature size and increase in J_c to 20 kA/cm², it is not difficult to envision that SCE chips may communicate across an MCM at 50 Gbps in the timeframe required. Further work must be done to address the area and power overhead of these driver circuits. Direct SFQ-to-SFQ signal transmission, which would greatly simplify chip floor planning and signal routing needs development.

2) *MCM-to-PCB Connection*: Three methods of high data rate motherboard-daughterboard connections have already been developed:

- A connector developed for the IBM superconductor computer project.
- Another connector developed by IBM for commercial applications using a dendritic interposer technology.
- A “beam-on-pad” approach developed by Siemens.

Although these three techniques are not immediately applicable due to either low interconnect density or low bandwidth, they serve as conceptual springboards for new designs. The main issue with the PCB to MCM interface is that approximately >500 Multi-Gbps lines must connect in the cross-sectional area of 50 mm², or 0.1 mm² per pin. For reference, a typical commercial “high-density,” low-speed connector has pins spaced on a 1.25 mm pitch for a density of 1.5 mm² per pin; a factor of 15 too large.

3) *MCM-to-MCM Connection*: The method of routing is to use short flexible ribbon cables, similar to those being used for the I/O cables. The density requirement of 2,000 per edge is considerably less than either for the external I/O cables or for the MCM-to-PCB. Given the length along the MCM bottom and top edges of 20 cm, the signal line pitch would be 4 mils, a density that is available today from flex circuit vendors.

While using short flex cables is the leading candidate technology to connect vertically adjacent MCMs, a second method of making these connections based on interposer technology was developed for another SCE program. This design has several advantages:

- Direct board-to-board communication without the use of a cable.
- The ability to “make and break” the connections many times.

This last aspect greatly increases the ability to assemble and repair the system.

3D Packaging

Conventional electronic circuits are designed and fabricated with a planar, monolithic approach in mind where there is only one major active device layer along the z-axis. Any other active layers in the third dimension are placed far away from the layers below and there are none (or only a few) connection between layers. The trend of placing more active devices per unit volume resulted in technologies referred to as 3D packaging, stacking, 3D integration. Compact packaging technologies can bring active devices closer to each other allowing short TOF, which is critical for systems with higher clock speeds. In systems with superconducting components, 3D packaging enables:

- Higher active component density per unit area.
- Smaller vacuum enclosures.
- Shorter distances between different sections of the system.

For example, 3D packaging will allow packing terabytes to petabytes of secondary memory in few cubic feet (as opposed to several hundred cubic feet) and much closer to the processor.

3D packaging was initiated in the late 70's as an effort to improve the packaging densities, lower system weight and volumes and improve electrical performance. Main application areas for 3D packaging are systems where volume and mass are critical ^{2,3} Historically, focal-plane arrays with on-board processing and solid-state data recorder applications for military and commercial satellites have driven the development of 3D packages for memories. For example, data storage unit of the Hubble Telescope consists of 3D stacked memory chips. Recently, 3D packaging has appeared in portable equipment for size savings. These applications have combined memory chips with controller or used few layers of stacked memory chips indicating that the cost of limited number of stacked layers is reduced to be acceptable to mainstream commercial applications. 3D packaging has expanded from stacking homogenous bare die (e.g., SRAM, Flash, DRAM) to stacking heterogeneous bare die and packaged die for “system-in-stack” approaches.

Two major types of stacking approaches are homogeneous and heterogeneous stacks. The first approach uses bare die as the layer of the stack for homogeneous stacking (meaning that each layer contains identical die). The I/O pads of the die are rerouted to the edge, and several die are laminated together to form the stack. The edge is lapped to expose the rerouted pads, and metalization is applied on one or more sides to interconnect layers and the cap substrate by connecting exposed pads. The intersection of the layer traces and the bus traces form the “T-Connects,” which are essential for high reliability. Unlike wrap-around or “L-Connect” approaches where step coverage problems can result in a reduced metal cross section, the trace thickness of T-Connects remains the full deposited thickness and contacts the entire cross-section of the exposed pad.

² S. Al-Sarawi, D. Abbott, P. Franzon, *IEEE Trans. Components, Packaging and Manufacturing Technology-Part B*, **21**(1) (1998) p. 1

³ V. Ozguz, J. Carson, in *SPIE Optoelectronic Critical Review on Heterogeneous Integration*, edited by E. Towe, (SPIE Press, 2000), p. 225.

For system-in-stack purposes, it was necessary to develop a technique to be able to stack packaged components and/or bare die components. This technique creates uniform size layers with non-uniform die sizes in each layer by creating a frame around the die (or dice) using potting as shown in Figure 4. Process starts with a bumped KGD encapsulated in a potting compound. Thin film metal traces contact the I/O bumps and route the signals to the edges of the potting material. The layer is ground from the backside resulting in a thin (100-250 micron) complete layer. The stack is created by laminating many layers with a cap substrate, which uses through-holes to provide I/O paths. Metalization is applied on one or more sides to interconnect layers and the cap substrate.

System-in-stacks with heterogeneous mix of components and the needed discretes for the high-speed operation of the stack can be realized. Multi-layer traces can be used to support complex interconnect structures needed for multi-chip layers and for high-speed operation up to 30 GHz. Heat management layers (such as copper, diamond, etc.) can be added on top of active components. These layers have dual functions:

- Providing a low-resistance, direct path for efficient thermal flow to the edge of the stack to keep the maximum temperature in the stack under control.
- Serving as shims to maintain the total layer-to-layer tolerances of the stack for high-density edge interconnect structures.

The manufacturability is maintained by allowing loose tolerances in the layer formation (component height variations, substrate thickness variations). Heterogeneous 3D chip stacking is a versatile technology allowing stacking of non-electronic components. Waveguides, vertical cavity surface emitting lasers, detectors are among devices used in stack layers. The free surfaces of the stacks are especially suitable for these types of devices. Optical interconnects from stack to stack and integrated focal plane arrays were already demonstrated⁴. Superconducting electronic circuits based on NbN and Nb Josephson junctions were inserted in stack layers and operated at temperatures as low as 4 K indicating the flexibility and reliability of the material system used in 3D chip stacks (Figure 5). Stacked systems were operated at 20 GHz. There are proposed approaches to integrate MEMS devices and passive optical components in stacks.

When the material selection and system design is judiciously performed, chip stacks have been shown to operate reliably with MTTF exceeding 10 years, in a wide temperature range (from – 270 C to 165 C) and in hostile environments subjected to 20,000 G. 3D packaging provides an excellent alternative to satisfy the needs high functionality system and sub-system integration applications, yielding system architectures that cannot be realized otherwise.

Thermal management is a key issue when aggressive miniaturization is used to pack large electronic functionalities into small volumes. When the power budget increases, thermal management layers can be inserted in the stack as alternating layers in addition to active layers. Experimental and analytical results indicated that up to 80W can be dissipated in a stack volume of 1cm³. Thermal resistance within the stack can be as low as 0.1 C/W. Minimal thermal resistance is critical for superconducting electronic applications where the components need to operate at 4 K.

Other trade-offs associated with 3D packaging are:

- Cost.
- Test.
- Reparability.

⁴ V. Ozguz, P. Marchand, Y. Liu, Proc. *International Conference on High Density Interconnect and System Packaging*, (2000), p. 8

We must note that the reparability is a common issue for each high-level integration approaches including VLSI. Testing of the complex 3D system-in-stack requires further development and better understanding of the system design. Added complexity will result from the high clock speed of SCE-based systems.

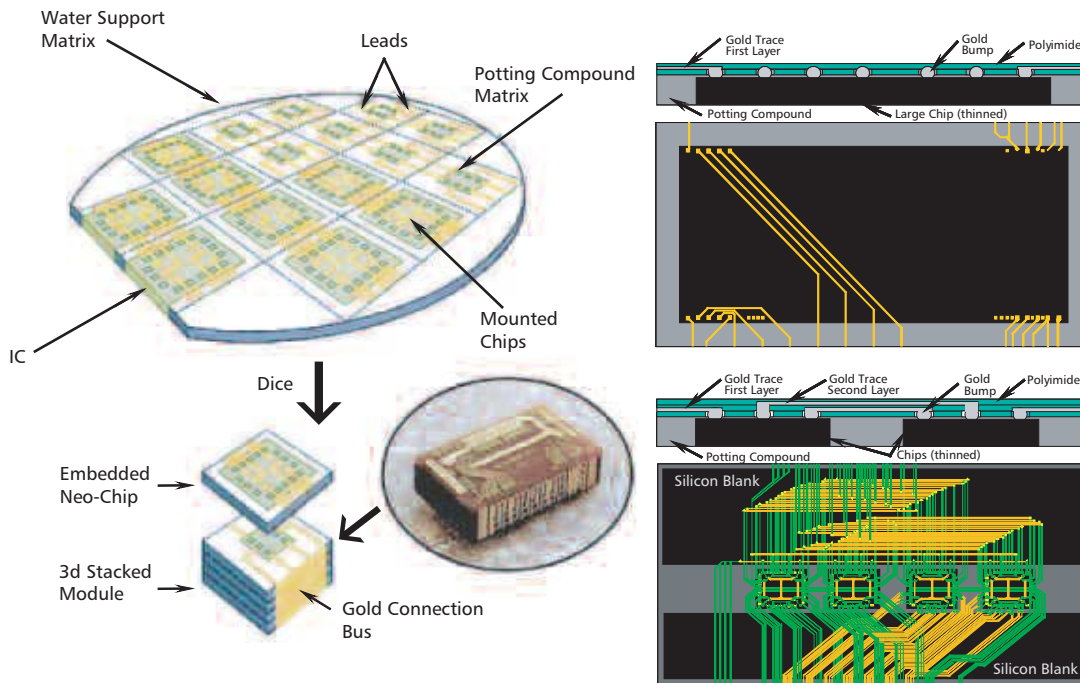


Figure 4. Heterogeneous stacking technology for system-in-stack.




Figure 5. Stacks containing superconducting electronics circuits were cycled from RT to 4 K several times; surviving 570 C temperature excursion. High speed operation was demonstrated up to 20 GHz in test layers.

Enclosures and Shields

Enclosures and shields are required to operate SCE:

- At cold temperatures.
- Under vacuum.
- Without magnetic interference, undue radiative heat transfer or vibration.

The vacuum dewar is an integral part of the cryocooler operation and should be designed concurrently. A hard vacuum is needed to avoid moisture formation in the system during the operation at low temperatures and to minimize convective heat transfer from the cold to warmer parts. Vibration of the system can translate into parasitic magnetic currents that may disrupt the operation of SCE circuits. Most of the engineering data such as magnetic noise levels, vibration levels, materials used and design techniques are trade secrets to various companies such as NGST, HYPRES and others who developed such systems in the past.

The basic operation of SCE circuits at temperatures as low as 4 K has been demonstrated by several companies as prototypes or as products. Most of these systems were single rack or single stage standalone systems. A few examples are shown in the following figures.

As can be seen in these figures, all of them use a circularly symmetrical packaging approach and are of a similar size, on the order of 20-100 cm. Other rack mounted designs can be envisioned for distributed systems.

Cryo packaging information about non-U.S. applications is scarce. Although several sources, such as NEC, report on their SCE circuit design and performance, they seldom report on their packaging approach. Reported test results may be derived from laboratory style apparatus.





Figure 6. A prototype cryo enclosure for operation at 4 K (Ref: NGST).



- I/O interface to the room temperature electronics adapted to the ribbon cable by developed hermetic ribbon cable connector
- DC wires vacuum interconnection is going to be done through weldable Deutsch connectors with ability of the differential bias input
- Vacuum can and top flange adapted for standard ISA flange
- Liquid helium damper implemented in the design



Cryogenic Module Dimensions
OD 11" X 25"

Figure 7. A typical cryocooler enclosure designed for communication applications (Ref: HYPRES).



Vacuum Can
with Magnetic
Shield



Cryogenic Module Dimensions
OD 11" X 25"



Figure 8. A rack mounted cryocooler enclosure example (Ref: HYPRES).

Several issues need to be resolved for working SCE systems. Modularity and assembly sequence are the most important issues for larger scale systems. Other major technical issues include the penetration of vacuum walls and shields with multitude of both low- and high-frequency cable, concurrent presence of high DC current and very high frequencies (50-100 GHz) digital signals and the thermal gradient from room temperature to 4 K and the resulting material compatibility requirements. The packaging is also a very strong function of the IO approach used to bring high data rate signals into and out of the SCE circuits. If optical data transfer is selected, the fibers and other OEO components represent an additional class of materials that may complicate system integration. Due to the lack of a previous design at large scale (e.g., a functional teraflop SCE supercomputer), other potential issues may be hidden and can only be discovered when such an engineering exercise takes place. The past funding (commercial and government) for the development of such systems was in disconnected increments that prevented the accumulation of engineering know-how. Therefore, substantial learning and more development is needed to insure a functional SCE-based supercomputer system.

Cooling

An SCE-circuit-based system needs to operate at temperatures ranging from 4 to 77 K. The cooling is performed by using cryocoolers or by liquid cooling using refrigerants such as Liquid Helium (LHe) or Liquid Nitrogen (LN₂). Liquid cooling:

- Often produces a smaller temperature gradient between the active devices and the heat sink than in vacuum cooling, making it easier to bias all the chips at their optimum design temperature.
- Also produces a more stable bias temperature due to the heat capacity of the liquid that tends to damp any temperature swings produced by the cooler, often allowing liquid cooled devices to perform better than in vacuum mounted parts.

For safety and practical reasons, liquid cooling is limited to 4 K (LHe) and 77 K (LN₂) regions. Liquid cooling which fails to recondense the evolved gas is not suitable for large scale devices, but closed systems with liquifiers are well established in the particle accelerator community where total cooling loads are large and direct emersion of parts into boiling cryogen is straightforward. Additionally, the boiled-off working fluid can be used to cool electrical leads going into the device, although this means they must be routed thru the vacuum vessel with the gas for some distance which may be unduly constraining. Small closed cycle coolers commonly use helium gas as a working fluid to cool cold plates to which the circuits are mounted in vacuum. By using mixed gases and other means, the temperatures of intermediate temperature stages can readily be adjusted. The most common type of small cryocoolers in the marketplace are the Gifford McMahon (GM) type cryocooler, but pulse tube coolers and Stirling cycle refrigerators which offer lower vibration and maintenance and better efficiency at a currently higher cost are also used.

In selecting the cooling approach for systems, the heat load at the lowest temperatures is a critical factor. A large-scale SCE-based computer with a single, central processing volume and cylindrically symmetric temperature gradient was analyzed under the HTMT program⁵. The heat load at 4 K arises from both the SCE circuits and the cabling to and from room temperature. The cable heat load (estimated to be kW) was the larger component, due to the large volume of data flow to the warm memory. Circuit heat load may be as small as several hundred watts. If all this heat at 4 K extracted via LHe immersion, a heat load of 1 kW would require a 1,400 liter/hour gas throughput rate.

Several implementation approaches that would divide the total system into smaller, separately cooled modules will be discussed in the issues section. Both an approach with a few (but still large) modules and an approach with many smaller modules have potential virtues when it comes to the ease of wiring and minimization of cable loads. However, there are availability differences between large scale and small scale cryo-coolers. Therefore, the status will be presented for both types separately⁶.

⁵ HTMT Program Phase III Final Report, 2002

⁶ "Integration of Cryogenic Cooling with Superconducting Computers", M. Green, Report for NSA panel.

Small Cryocoolers

Among commercial small cryocoolers, the GM coolers have been around the longest; tens of thousands of GM units have been sold. Single stage coolers that go down to 25 to 30 K have been around since 1960. Two stage coolers that go down to 7 K have been used since the mid 1960's and 10 K versions are extensively used as cryopumps (condensation) to keep Si process equipment at low pressure. With the use of newer regenerator materials, two stage GM coolers now can go down to temperatures as low as 2.5 K. GM coolers have a moving piston in the cold head and as a result, produce vibration accelerations in 0.1 g range. GM coolers are quite reliable if scheduled maintenance of both the compressor units and the cold heads is performed regularly. A typical maintenance interval is about 5,000 hours for the compressor and about 10,000 hours for the cold head.

In the hybrid cooler and 77 K communities, there are a number of companies selling Stirling cycle coolers (e.g Sunpower and Thales) but there are no commercial 4 K purely Stirling coolers on the market. Stirling cycle coolers with cycle efficiency of up to a factor of two have been demonstrated, but whether this will prove a general advantage for Stirling machines is unclear. The units tend to be physically smaller than GM machines of similar lift capacity due to the absence of oil separation system. The vibration characteristics of Stirling machines are about the same as for GM machines. Free piston machines with electronically controlled linear compressors and expansion engines have been made with substantially lower vibration for applications such as on-orbit focal planes. At this time, and under conditions where maintenance is feasible, commercial Stirling machines do not appear to be more reliable than GM machines.

Pulse tube coolers are commercially available from a number of vendors including Cryomech, Sumitomo, and Thales. Two stage pulse tube coolers can produce over 1 W of cooling at 4.2 K on the second stage of a two stage cooler and simultaneously produce about 50-80 W at 50 K. Pulse tube cooler have gone down to 2.5 K (1.5 K with He3). However, there is a conflict in designing a single compressor pulse tube between the high pump frequency which produces good efficiency at the higher temperatures and the low frequencies needed at low temperatures. Hybrid Sterling-pulse tube coolers allow the higher efficiency of a Sterling high-temperature machine to be mated to an efficient low frequency pulse tube bottom end, at a cost of having two compressors. A pulse tube machine will inherently produce lower vibration accelerations (<0.003 g) on the cold head than a GM cooler. Since there are no moving parts in the cold head, pulse tube machine are easier to make reliable. Maintenance intervals of 25,000 hours are recommended for the Cryomech machines. Pulse tube cryocoolers were favored for space applications due to their low vibration and higher reliability without maintenance when flexure bearing compressors are also used. There are many companies involved in creating space cryocooler technology including Lockheed-Martin, Ball Aerospace, Creare, ETA Incorporated, Honeywell, Hughes, Mainstream Engineering Corporation, Mitchell/Stirling Machine Systems, Northrop Grumman (including TRW), Raytheon, Ricor, and Swales Incorporated. The development was funded by DoD components, and the results are not commercially available.



Figure 8. Northrop Grumman's high-efficiency cryocooler unit weighs less than 7 kilograms compared to previous state-of-the-art at 22 kilograms. Sixteen units have been ordered for various flight programs.

The Northrop Grumman Space Technology (NGST) single-stage High Efficiency Cryocooler (HEC) program, quite possibly the most successful to date, was a joint effort funded by MDA, NASA, and the U.S. Air Force (Figure 8). A 35/85 K NGST High-Efficiency Pulse Tube Cryocooler is intended to simultaneously cool both focal planes and optics, removing 2 watts at 35 K and 16 watts at 85 K with a single cooler. A parallel program with Raytheon has similar objectives but uses a hybrid Stirling upper stage and pulse-tube lower stage for lower mass and improved efficiency.

Large Cooling Units

There are two European manufacturers of large machines: Linde in Germany and Air Liquid in France. There are a few Japanese manufacturers, but their track record is not as good as the Europeans. Dresden University in Germany is one of the few places in the world doing research on efficient refrigeration cycles. A system capable of handling kW level of heat loads cools portions of the Tevatron accelerator at Fermilab in Illinois. Systems of such size produced a refrigeration efficiency of 0.25 (Figure 9). The conceptual design of a single large system has been initiated under the HTMT program (Figure 10).




Figure 9. Refrigerator efficiency of various cooling systems.




Figure 10. Concept for a large-scale system including cryogenic cooling unit for supercomputers⁷.

⁷ HTMT Program Phase III Final Report, 2002

Although cooling of the 4 K circuits by immersion in LHe may be the most straightforward approach cryogenically, doing so would produce significant technical issues on the system integration front: large diameter plumbing would be required to carry the helium gas back to be recondensed and utilizing the enthalpy of the gas in order to pre-cool the I/O lines requires a complicated heat exchanger. Most importantly, the 4 K components would need to be surrounded by a He leak-tight vessel and then a vacuum space. Either the leads would have to travel to elevated temperatures with the gas, necessitating long signal travel times, or they would travel a shortest distance path that would require cryogenic vacuum feed-throughs for each of the leads. It seems likely that this approach practical only if substantial cold memory exists in the system so that only computational results are shipped to the room temperature. Otherwise, millions of high speed I/Os are needed

Cooling by conduction using plates made from high thermal conductivity materials such as copper seems a more viable alternative when SCE chips, MCMs, PCBs, and I/Os are all in the same vacuum space. These plates would either have direct thermal connections to a cold plate, cooled by the refrigerator, or have liquid helium running through a set of engineered sealed passages (e.g. micro-channels or pores of Cu foam) within the plates themselves. Such systems have recently been extensively worked to provide extreme rates of heat removal from above room temperature for power amplifiers and the same physics will apply. Whether both the liquid and gas phase of He can be present might need to be investigated. Conduction cooling places greater demands on the thermal design of the system but the technical issues are familiar. Low thermal conductance paths from the chip to the cold head are essential. The temperature difference between the component being cooled and the cold box is a key issue even when large refrigerators are used to cool superconducting electronic devices. Given the severe energy penalty for compensating for this gradient by lowering the cold-plate temperature, bit-error tests should be performed for 20 KA/cm² circuits as a function of bias temperature as early as possible. For 1KA/cm² devices there is experimental evidence that circuits designed for 4.2 K operation still function well as high as 5.5 K due to the slow dependence of critical current (I_c) on temperature (ref. HYPRES).

Serious consideration should be given toward reducing the refrigeration load through the use of high temperature superconducting (YBCO or MgB₂) cables or optical leads in conjunction with mid temperature (e.g. 30- 40 K) intercepts. More details are provided in the "Cables" section.

Vibration can be a serious problem for superconducting electronics if the total magnetic field at the circuits is not uniform and reproducible. Good cryo-packaging practice requires both careful attention to magnetic shielding, minimization of vibration and relative motion of circuit components. If it proves essential, the sensitive electronic components can be mounted using a stiff (high resonant frequency) cold mass support system that has a low heat leak and is cooled via a flexible cooling strap attached to the cold plate. Keeping the strap short length minimizes the temperature drop and also reduces the vibration isolation achieved. The cryocooler itself can be mounted with flex mounts on the cryostat vacuum vessel, which has more mass than the cooled device.

The reliability issues can be mitigated by using smaller cryocooler units around a central large scale cryogenic cooling unit. This approach adds modularity to the system and allows for local repair and maintenance while keeping the system running. However, the design issues resulting from the management of many cryocoolers working together are yet not well explored.

Another issue is the cost of the refrigeration. Buildings and infrastructure may be a major cost factor for large amounts of refrigeration. Effort should be made to reduce the refrigeration at 4 K and at higher temperatures since the cost (infrastructure and ownership) is direct function of the amount of heat to be removed (Figure 11).




Figure 11. The cost of various refrigerators as a function of the refrigeration.

One key issue is the availability of manufacturers. The largest manufacturer of 4 K coolers is Sumitomo in Japan. Sumitomo has bought out all of its competitors except for Cryomech. There No American company has made large helium refrigerator or liquefier in the last 10 years, the industrial capacity to manufacture large helium plants having died with the Superconducting-Super-Collider in 1993. Development funding may be needed for U.S. companies to insure that reliable American coolers will be available in the future. Development toward a 10 W or larger 4 K cooler would be desirable to enable a supercomputer with a modular cryogenic unit.

Power Distribution and Cables

SFQ circuits are DC powered. Due to the low voltage (mV level), the amount of current to be supplied is in the range of few Amps for small scale systems and can be easily kiloAmps for large scale systems. Therefore, special precaution needs to be taken to provide DC power to the system.

Large currents require low DC resistances that translate to large of amount of metal lines. However lower thermal conductance implies high resistance cables, since the thermal conductance of a metallic line is inversely proportional to the electrical resistance. High electrical resistance gives high Joule heating for dc bias lines and high attenuation for ac I/O lines. Therefore, each I/O design must be customized to find the right balance between thermal and electrical properties. Reliable cable attach for thousands of connection requires also cost efficient assembly procedures with high yield.

Supplying DC current to all of the SCE chips in parallel would result in a total current of many kiloAmperes. Several methods may be used to reduce this total current and heat load. One technique is to supply DC current to the SCE chips in series, rather than in parallel (current recycling) taking advantage of very low resistances of lines in SCE circuits. However, it would require true differential signal propagation across floating ground planes. High temperature superconductor (HTS) cables may be used to bring in DC current from 77 K to the 4 K environment. Another solution is to use switching power supplies: high voltages/low currents can be brought near SCE circuits and conversion to low voltages/high currents, all at DC, can occur at the point of use. However, this method employs high power field-effect transistor switches, which themselves can dissipate significant power.

Of these methods, current recycling shows the most near-term promise. Laboratory level experiments were conducted for small scale systems. The demonstration of a large-scale SCE system with kiloAmperes of current was not performed yet. For smaller IO counts, coaxial cables can be used for both high and medium-speed lines. Coax cables have different sections. The middle section is a short length of stainless steel coaxial cable that has a high thermal resistance and a bottom section of non-magnetic Cu coaxial cable for penetration into the magnetically-shielded chip housing. For higher-frequency input signal up to 60 GHz and clock lines, Gilbert-Corning GPPO connectors are used. Flexible BeCu microstrip ribbon cables were also considered for medium speed output lines. At 1 GHz, electrical attenuation and heat conduction of these ribbon cables were measured to be nearly identical to the tri-section coax cable. Modified connector assemblies were tested at room temperature and at liquid nitrogen temperature and found to be within specification for both reflection and attenuation.

For systems with hundreds or thousands of I/O lines, coaxial cables are not practical. To meet this challenge flexible ribbon cables have been developed⁸. These cables consist of two or three layers of copper metallization separated by dielectric films, typically polyimide. With three copper layers, the outer two layers serve as ground planes and the inner layer forms the signal lines, creating a stripline configuration. Stripline cables provide excellent shielding for the signal lines. The dielectric layers consisted of 2 mil thick polyimide films. The ground planes were fabricated from 0.6 micron-thick sputtered-copper films, while the signal lines were patterned from a 4-micron-thick plated-copper film. The signal line width was 3 mils for 50 Ω impedance with a pitch of 70 mils and each cable has more than hundred leads. In addition to large I/O count and multi-GHz connections, flexible ribbon cables offer the advantage of low thermal conduction to the cold stage. Cables can be manufactured with copper films the thickness of the electrical skin depth for a particular frequency of interest. Successful cable-attach procedures with high reliability and low cost were also demonstrated. Signal speeds up to 3 GHz were experimentally demonstrated. The construction of the cable should allow operation up to 20GHz after minor modifications.

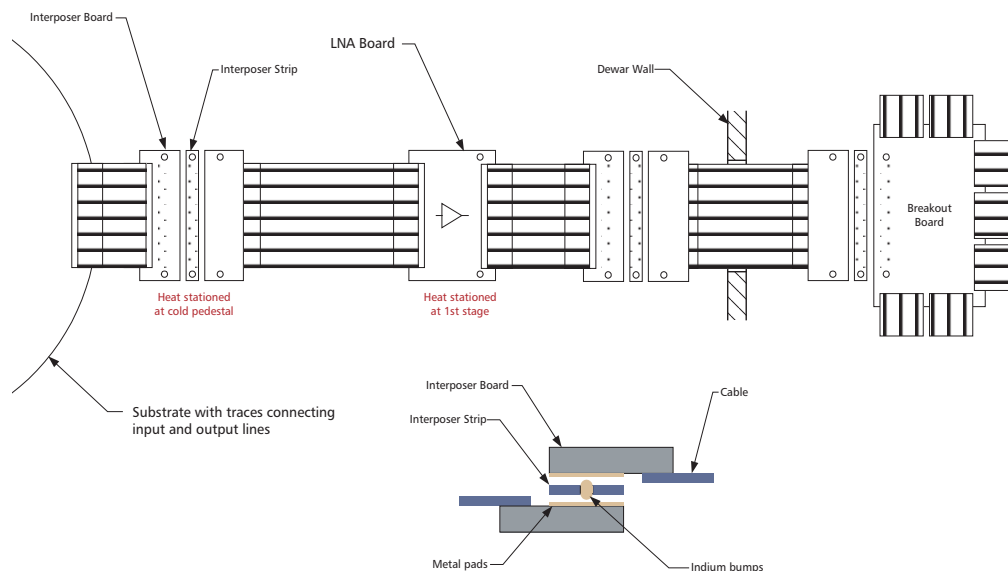


Figure 12. A high-speed flexible ribbon cable designed for modular attachment (Ref: NGST).

⁸ "Cryogenic Packaging for Multi-GHz Electronics" T. Tighe et al, IEEE Tran. Applied Superconductivity, Vol 9 (2), pp3173-3176, 1999.

To maximize both the magnetic and radiation shielding in digital SCE systems, any gaps in the shields, such as those required to feed in I/O cables, need to be minimized. A significant benefit of flexible ribbon cables are their ~0.1" bending radius and <10 mil total thickness. These cables easily snake in and out of the shields, and their thinness allows for exceedingly narrow apertures to be used as feed-throughs. In addition, the wide ground planes allow for convenient and effective places to attach heat sinks, thus further reducing the heat load to the cold stage.

Use of fiber optics to carry information into the cryostat reduces parts count and thermal load. Using wavelength division multiplexing (WDM), each fiber can replace 100 wires. Optical fibers have lower optimal thermal conductivity than metal wires of the same signal capacity. A factor of 1,000 reduction in direct thermal load is achieved, as each fiber has 10 times less thermal conductance relative to wire cable. Using optical means to bring data *into* the cryostat has been sufficiently demonstrated to be viewed as low risk. These advantages apply only to data traveling into the cryostat. There are known low-power techniques to convert photons to electrical current, but the reverse is not true. Small signal strength of SFQ complicates optical signal output. Magneto-optical coupling of the signals at an intermediate temperature appears the most viable scenario among the ideas already explored. Driving optical signals at high frequencies requires lasers or diodes. An average laser channel requires 1 mA of current at 1 V for a power consumption of 1 mW. This compares unfavorably with the present estimate of 50 μ W thermal loads for each electrical signal line. At the same time, direct modulation of lasers or diodes at 4 K is just one method of transmitting optical data. Another method is to generate the photons at room temperature and modulate the amplitude of the optical signal at 4 K. In theory, this promises greatly reduced power consumption within the cryostat. At present, however, a low-power modulation technique at multi-Gbps data rates is unproven. More details are presented in the "Interconnects" section.

An alternate material to metals for the electrical wires is HTS. HTS cables could be used between 4 K and 77 K. Conductors made from HTS materials, in theory, transmit electrical signals with little attenuation and with little heat load. This combination—which cannot be matched by non-superconductor metals—makes them attractive. However, the cuprate "high" T_c materials are inherently brittle in bulk form and so lack the flexibility that is important to the assembly of systems with the 1,000's of leads expected in a supercomputer. The second generation magnet wire products based on YBCO thin films on a strong and flexible backing tape now coming into commercial production are being modified in an ONR phase 2 award for use as flexible DC leads. Carrying tens to hundreds of A from ca 77 K down to 4 K with low thermal loading should be quite straightforward.

The same product is inappropriate for RF or digital signals due to the high loss expected from the metal in the support tapes. A prototype coplanar RF line product was demonstrated < 2,000, but was never commercialized and would have inadequate packing density for use in a supercomputer.

Given the large material difficulties of the cuprates, a better choice for the RF leads would be MgB_2 which was recognized to be a 39K, s-wave superconductor only late in 2001. However, it has attracted considerable attention in the filter community due to its much simpler physics and better in-B field properties than the cuprates and much higher transition temperature than the traditional LTS materials. A scalable manufacturing process has been demonstrated at Superconductor Technologies that could today manufacture 10 1 cm x 5 cm tapes using the current substrate heater and could be straight-forwardly scaled up to a process that simultaneously deposits over a 10-inch diameter area or onto a long continuous tape. Deposits on both sides of carrier wafer/tape are feasible, and any substrate for which there is not a chemical reaction can be used. The films are of very high quality, regularly having dc resistivity values and microwave loss tangents (about $2 \times 10^{-5} \Omega$ at 10 K and 10 GHz) some of the lowest reported. Indeed, this loss tangent is lower than both YBCO and Nb at low temperatures. At 30 K, the thermal conductivity is ~ 0.3 W/cm-K. Ignoring the expected drop in this conductivity at lower temperatures, a 1 cm wide tape with 0.5 microns of MgB_2 and 5 cm length to span 30 to 5 K would conduct 0.075 mW through the MgB_2 . The films are relatively stable with respect to patterning.

These properties enable the possibility of making flexible MgB_2 thin-film leads for high-speed signal propagation. The next step in a development effort could include demonstration of MgB_2 film growth on suitable flexible substrate tapes of the size discussed above. By making the tape from Si_3N_4 over an etched away grid of Si screening, the thermal contribution of the supporting tape could be minimized. Alternatively, growth on single-crystal YSZ and polycrystalline alumina substrates without the use of a buffer layer have already been achieved. Thus growth on flexible YSZ tape may be straightforward. Growth on other materials with suitable high-frequency properties (such as the flexible “RT/duroid”) should also be investigated. The two sided growth allows a strip line geometry with the substrate as the dielectric to be realized, but work on deposited dielectrics is desirable to minimize signal cross talk and increase signal trace density. It is estimated that such tapes could be productized within four years if an engineering development program is implemented.

The effects of the current recycling on the high frequency operation of SCE circuits require further development and testing. The operation at 50 GHz and beyond with acceptable bit-error-rates (BER) needs further development and testing, including physical issues (e.g., conductive and dielectric material selection, dielectric losses) and electrical issues (signaling types, drivers and receivers).

The combination of good electrical and poor thermal conductance is inherently difficult to meet since a good electrical conductor is also a good thermal conductor. Specialized coaxial cables have been commonly used for multi-GHz I/O connections and can support operation up to 100 GHz. These cables employ relatively low thermal conductivity materials such as silver-plated steel and beryllium-copper, and can have diameters down to 20 mils. However, hand-wired coax is not practical for systems that require thousands or millions of I/O lines.

The required density of I/O lines in large scale systems is a factor of 10 greater than the density achievable today (0.17 lines/mil) for flexible ribbon cables. One method to increase the apparent density of the cables is to piggy back or stack several flex cables, but this does lead to increased complexity and may limit the MCM-to-MCM spacing.

Optical interconnects may requires the use of elements such as prisms or gratings. These elements can be relatively large, and may not fit within the planned size of the cryostat. The power required for the optical receivers, or the thermal energy delivered by the photons themselves may offset any gains from using the low thermal conductance glass fibers. A detailed design is needed to establish these trade-offs.

The use of HTS wires includes two significant technical risks:

- Currently, the best HTS films are deposited epitaxially onto substrates such as lanthanum-aluminate or magnesium oxide. These substrates are brittle and have relatively large thermal conductance which offset the “zero” thermal conductivity of the HTS conductors.
- Any HTS cable would need to have at least two superconductor layers (a ground plane and a signal line layer) in order to carry multi-Gbps data.

Presently, multilayer HTS circuits can only be made on a relatively small scale due to pin-holes and other such defects which short circuit the two layers together.

For additional copies, please contact:

NSA Office of Corporate Assessments
9800 Savage Road, Suite 6468
Fort George G. Meade, Maryland 20755

Attention: Chief, Technology Test and Evaluation